

Kim, Jeongbin; Lim, Wooyoung; Schweighofer-Kodritsch, Sebastian

Working Paper

Bargaining and Time Preferences: An Experimental Study

CESifo Working Paper, No. 8683

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Kim, Jeongbin; Lim, Wooyoung; Schweighofer-Kodritsch, Sebastian (2020) : Bargaining and Time Preferences: An Experimental Study, CESifo Working Paper, No. 8683, Center for Economic Studies and Ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/229501>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

**Bargaining and Time
Preferences:
An Experimental Study**

Jeongbin Kim, Wooyoung Lim, Sebastian Schweighofer-Kodritsch

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Bargaining and Time Preferences: An Experimental Study

Abstract

We generalize the [Rubinstein \(1982\)](#) bargaining model by disentangling payoff delay from bargaining delay. We show that our extension is isomorphic to generalized discounting with dynamic consistency and characterize the unique equilibrium. Using a novel experimental design to control for various confounds, we then test comparative statics predictions with respect to time discounting. All bargaining takes place within a single experimental session, so bargaining delay is negligible and dynamic consistency holds by design, while payoff delay per disagreement round is significant and randomized transparently at the individual level (week/month, with/without front-end delay). In contrast to prior experiments, we obtain strong behavioral support for the basic predictions that hold regardless of the details of discounting. Testing differential predictions of different forms of discounting, we strongly reject exponential discounting in favor of present-biased discounting.

JEL-Codes: C780, C910, D030.

Keywords: alternating-offers bargaining, time preferences, present bias, laboratory experiments.

Jeongbin Kim
National University of Singapore
Centre for Behavioural Economics (CBE)
Singapore 117602
bizkj@nus.edu.sg

Wooyoung Lim
The Hong Kong University of Science and
Technology, Department of Economics
Kowloon / Hong Kong
wooyoung@ust.hk

Sebastian Schweighofer-Kodritsch
Humboldt-University Berlin
School of Business and Economics
Germany – 10178 Berlin
sebastian.kodritsch@hu-berlin.de

September 7, 2020

We are grateful to Andreas Blume, Anujit Chakraborty, Dongkyu Chang, Pedro Dal Bó, Michael Böhm, John Duffy, Yoram Halevy, Emin Karagözoğlu, Shih En Lu, Johannes Maier, Philip Neary, Thomas Palfrey, Ariel Rubinstein, and Charles Sprenger for valuable comments and suggestions. For helpful comments and discussions, we thank seminar participants at Caltech, HKUST, Hong Kong Baptist University, KAIST, National Taiwan University, Peking University Guanghua School, University of Hong Kong, University of Vienna, University of Bonn (briq), and conference/workshop participants at the 2019 ESA World Meeting at Simon Fraser University, the 2019 CPMD Workshop on Behavioural and Experimental Economics at University of Technology Sydney, the 2018 HKUST Workshop on Experimental Economics, and the 2019 SUFE Workshop of Behavioral and Experimental Economics. This study is supported by a grant from the Research Grants Council of Hong Kong (Grant No. GRF-16500318), and Sebastian Schweighofer-Kodritsch also gratefully acknowledges financial support from the Deutsche Forschungsgemeinschaft (through CRC TRR 190).

1 Introduction

How will two parties to a transaction divide the economic surplus that it creates? As fundamental as this question is to economics, it was not until the work of [Rubinstein \(1982\)](#) that economic theory offered a useful answer in terms of a tractable strategic model with a clear prediction based on the parties’ individual preferences. His disciplined model of the bargaining process as one of indefinitely alternating offers constitutes the core of modern non-cooperative bargaining theory; it has been extended in many directions and applied to various settings.¹

The central driving force of this theory is the parties’ individual costs of the *payoff delay* that results whenever they disagree. Their time preferences determine their relative bargaining power. Not surprisingly, given the model’s importance, several studies have experimentally tested its predictions. However, the costs of disagreement in all of these studies have been implemented either as a reduction in monetary surplus or as a risk of exogenous breakdown rather than as a payoff delay. Hence, whether actual time preferences behaviorally affect the bargaining outcome as predicted has remained an open question.

In this paper, we introduce a novel experimental design that directly addresses this question. In contrast to the prior experimental literature, we find strong behavioral support for all fundamental predictions of the theory once we account for present bias. Quite remarkably, this is in line with the large body of empirical studies that directly measure time preferences and document present bias as the most important qualitative deviation from exponential discounting (e.g., [Frederick, Loewenstein, and O’Donoghue, 2002](#); [Augenblick, Niederle, and Sprenger, 2015](#); [O’Donoghue and Rabin, 2015](#)).

Our key innovation, both theoretically and experimentally, is to disentangle the timing of payoffs from the timing of bargaining rounds. Theoretically, we generalize the classic [Rubinstein \(1982\)](#) model to arbitrary payoff delays per round of disagreement and general time preferences, with the only substantial assumption that preferences are dynamically consistent. We establish and characterize the unique (subgame perfect Nash) equilibrium of this game, which involves history-independent strategies and has immediate agreement in every round. When the delay between any two rounds of bargaining is negligible—i.e., offers are frequent, so that bargaining is essentially instantaneous—preferences are dynamically consistent in this game by design. Our theory therefore also covers any “naturally” dynamically inconsistent time preferences, such as quasi-hyperbolic and hyperbolic discounting, but without confronting any issues of *intra*-personal conflict or naïveté that arise when, instead, decisions are made over a significant time horizon.

The game we experimentally implement exploits this theoretical observation: All bargaining takes place within a single session, with frequent offers, whereas actual payoffs are subject to significant

¹The books by [Osborne and Rubinstein \(1990a\)](#) and [Muthoo \(1999\)](#) provide overviews. Relatedly, [Binmore, Rubinstein, and Wolinsky \(1986\)](#) show how the model provides a non-cooperative foundation to the Nash bargaining solution ([Nash, 1950](#)), which has been a major modeling device used in applied theory and empirical work (for a recent example, see [Ho and Lee, 2017](#)).

delay per round of disagreement.² Thus, we avoid the major concern with actual longitudinal designs (e.g., [Sprenger, 2015](#)) that attrition may be systematically related to time preferences.³

While the theory parsimoniously focuses on time preferences and assumes these are common knowledge, people’s preferences also include other relevant concerns—in particular, fairness and risk—and are highly heterogeneous. Strategic interaction is therefore inherently subject to incomplete information (relatedly, see [Fanning and Kloosterman, 2019](#)). We employ a novel experimental approach, recently introduced by [Kim \(2019b\)](#) as the *effective discounting procedure*,⁴ to control for various preference “confounds” and establish a setting in which time preferences nonetheless translate into clear comparative statics predictions to be tested between subjects. The procedure randomly assigns participants their individual length of payoff delay—either a week or a month per round of disagreement, and either including a front-end delay or not—and thereby creates groups of bargainers that are similar in terms of their *underlying general preferences* but differ in their *effective time preferences*. The payoff delay profiles that bargainers individually face, their “types,” are made common knowledge within any bargaining match.

Our experimental treatments correspond to particular matchings of such bargainer types. We implement three of these: Treatment *WM* matches bargainers whose payoff gets delayed by one week per round of disagreement (“weekly bargainers”) with bargainers for whom this is one month (“monthly bargainers”); Treatment *WM2D* is similar, except that every bargainer’s payoff comes with an additional front-end delay of one week (even in case of immediate agreement, and we call these bargainers “delayed”); Treatment *WW1D* matches two weekly bargainers of whom exactly one faces such a front-end delay. Within each treatment, participants play ten games under random rematching (subject to the treatment condition), with their individual type fixed throughout. In every such game, which type gets to make the initial proposal is determined randomly, so we observe both versions of the alternating-offers game for any type match in a treatment.

Under symmetry in terms of *underlying* preferences, assumptions on time preferences translate our type manipulation into differences in bargaining power. These comparative statics predictions in time preferences are our main focus, and we test them by comparing the cumulative distribution functions of behavioral measures of interest for first-order stochastic dominance. Although conservative, we deem this approach appropriate because our manipulation is supposed to effectively “shift” the entire distribution of time preferences.

The natural benchmark is constant/exponential discounting (EXD), for which our model reduces

²Our study is thereby tightly linked to the experimental evidence on time preferences elicited from binary choices between various time-dated monetary rewards. Given the indefinite strategic interaction, to ensure the credibility of our experiment, we additionally impose a commonly known 25% chance of random exogenous termination. This probability is held constant across all rounds of all games in all treatments. Strictly speaking, time preferences should be interpreted as including this risk here.

³[Kim \(2019a\)](#) indeed finds that patience and present bias measured at the beginning of his experiment closely predicted how long participants would take part in his longitudinal study.

⁴We thank John Duffy for helping coin this term.

exactly to that of [Rubinstein \(1982\)](#). Under EXD, the first two treatments are equivalent, and the third treatment is symmetric because a front-end delay is strategically irrelevant, akin to a sunk cost. In addition to the general predictions of immediate agreement and a proposer advantage, EXD clearly implies that weekly bargainers obtain a greater share than monthly bargainers in the same initial role.⁵ Our data from Treatments *WM* and *WM2D* strongly confirm this *basic delay advantage*.⁶ However, in contradiction to EXD, we find that the two treatments are not equivalent behaviorally, and that the behavior in Treatment *WW1D* violates symmetry.

The leading alternative to EXD is quasi-hyperbolic discounting (QHD), which adds a single parameter to capture the empirically well-documented phenomenon of a present bias ([Phelps and Pollak, 1968](#); [Laibson, 1997](#)). This model’s key implication is that a front-end delay is strategically advantageous to initial respondents.⁷ Relative to EXD, QHD therefore clearly breaks the symmetry in Treatment *WW1D* in favor of the delayed (weekly) bargainers. Our data from this treatment indeed strongly confirm this *front-end delay advantage*. Similarly, QHD also breaks the equivalence between Treatments *WM* and *WM2D* under EXD, in the direction of an advantage of respondents under the latter treatment, where they are delayed. While we find some support for this prediction when we compare those games where the initial proposer is a weekly bargainer, the opposite is true when we compare those games where the initial proposer is a monthly bargainer.

This latter finding, in turn, is consistent with hyperbolic discounting, which also features a present bias but as an instance of uniformly diminishing impatience. QHD can be interpreted as a parsimonious approximation that focuses solely on present bias (see [Frederick et al., 2002](#)). Adding a front-end delay shifts all delays into the future, and under hyperbolic discounting, this not only “removes” any present bias but effectively increases patience uniformly. In relative terms, this may increase bargaining power more for a monthly bargainer than a weekly bargainer. Indeed, a pronounced present bias together with hyperbolic discounting of future delays yields exactly those clear predictions from QHD that our data strongly confirm, while at the same time rationalizing the behavioral findings that contradict QHD’s predictions.⁸ Considering that hyperbolic discounting and present bias are exactly the key qualitative properties of empirically measured time preferences, across a huge number of studies eliciting time preferences from *individual* decisions, we interpret our findings as strong confirmation of the theory.

Overall, we conclude that time preferences are certainly not all that matters in bargaining, but

⁵We find a very high degree of efficiency in our experimental bargaining games, as almost three quarters end in immediate agreement, without significant differences across treatments. We also confirm a proposer advantage.

⁶We also strongly confirm this prediction in a comparison of behavior across Treatments *WM* and *WW1D*, which have the weekly bargainers (with no delay) as a common type. The other comparison across treatments for this prediction, namely *WM2D* and *WW1D*, which have the delayed weekly bargainers as a common type, violates it, but see below for present bias and hyperbolic discounting.

⁷Note that from Round 2 onwards all agreements have delayed payoffs, so present bias ceases to matter, and the initial proposer’s strategic advantage means that only the respondent’s bias materializes in terms of equilibrium.

⁸To obtain clear predictions, we impose minimal structure on hyperbolic discounting satisfied by all commonly considered models of such discounting; see [Loewenstein and Prelec \(1992\)](#) for a general parametric family. Closely related are [Halevy \(2008\)](#) and [Chakraborty, Halevy, and Saito \(2020\)](#), who relate present bias and hyperbolic discounting to the inherent uncertainty about future consumption and common violations of expected utility.

they do matter significantly and in a manner that is theoretically predicted by and consistent with what we know from the large body of work that has researched them.

Especially with respect to the negative conclusions from prior experimental investigations of the [Rubinstein \(1982\)](#) model (to be discussed in the next section), our findings may therefore be regarded as an unprecedented behavioral success story for the basic model of non-cooperative bargaining theory when extended to incorporate the key properties of real time preferences. With the latter qualification, they lend encouraging behavioral support to the large and important literature applying this model. Moreover, our finding that people seem to commonly understand and strategically respond to present bias not only supports the presumed prevalence of this bias but also promotes (further) theoretical work on dynamic strategic interaction with present-biased individuals. Finally, viewed from a somewhat different perspective, our results may also be taken as a demonstration of the methodological value of the effective discounting procedure for further experimental work on the role of time preferences in strategic settings.

The rest of this paper is organized as follows. The next section discusses the most closely related literature. We then present the general theoretical background for our study in [Section 3](#). This is followed by our experimental design, the behavioral predictions for the most important classes of time preferences, and administrative details in [Section 4](#). We report and discuss our experimental findings in the subsequent [Section 5](#). Finally, [Section 6](#) offers a brief conclusion. All proofs are relegated to [Appendix A](#). [Appendices B](#) and [C](#) provide additional figures and results on learning, and [Appendices D](#) and [E](#) contain experimental instructions (for one exemplary experimental treatment) and selected screenshots. [Appendix F](#) presents details of an additional time preference elicitation and results on how these measures relate to bargaining behavior, which support the basic approach of our study.

2 Literature Review

Our review of the literature focuses on (1) theoretical analyses of time preferences in the canonical bargaining environment with an infinite horizon and alternating offers and (2) experimental studies that investigate this bargaining model. There are large areas of work on bargaining that we do not cover, including the vast experimental literature on ultimatum bargaining and finite-horizon sequential bargaining, and the theoretical literature that extends the original [Rubinstein \(1982\)](#) model in several other directions, such as multilateral bargaining, bargaining with asymmetric/incomplete information, and endogenous proposer determination. For a recent review of the ultimatum bargaining literature, see [Güth and Kocher \(2014\)](#). For a comprehensive survey of non-cooperative bargaining theory during its most active period of research, see [Binmore, Osborne, and Rubinstein \(1992\)](#); for a more recent survey focusing on incomplete information, see [Ausubel, Cramton, and Deneckere \(2002\)](#).

Theory. In his seminal paper, [Rubinstein \(1982\)](#) introduces the canonical bargaining model in which two players alternate in making offers to each other on how to divide a given surplus until they reach

agreement. Assuming exponentially discounted concave utility and perfect information, there is a unique subgame-perfect Nash equilibrium. This equilibrium occurs in stationary strategies that imply immediate agreement in any round, hence efficiency. Given impatience and that the burden of delay is with the player responding to an offer, a proposing player enjoys a strategic advantage. Moreover, *ceteris paribus*, the more patient a player is—in particular, the higher her discount factor for given utility—the greater her bargaining power in the sense of capturing a larger share of the surplus in the equilibrium agreement. With symmetric preferences, as offers become infinitely frequent and players approach perfect patience, the proposer advantage vanishes, and the equilibrium outcome converges to an immediate equal split, as prescribed by the [Nash \(1950\)](#) bargaining solution.

Motivated by empirical evidence, several theoretical attempts have recently been made to generalize this model in terms of time preferences. Almost all have focused on “stable” preferences to maintain the game’s stationarity property, which makes the game tractable. In this case, any deviation from exponential discounting implies dynamic inconsistency, and [Schweighofer-Kodritsch \(2018\)](#) provides a comprehensive equilibrium characterization under minimal preference assumptions when these preferences are common knowledge, implying “full sophistication” (for related work see also [Ok and Masatlioglu, 2007](#); [Noor, 2011](#); [Pan, Webb, and Zank, 2015](#); [Lu, 2016](#)). He finds that with concave utility, a weak present bias is sufficient for a unique equilibrium similar to exponential discounting. However, as [Akin \(2007\)](#) and [Haan and Hauck \(2019\)](#) show for quasi-hyperbolic discounting, naïveté about present bias may lead to even perpetual disagreement.

We provide a hitherto overlooked alternative but formally equivalent interpretation to [Rubinstein’s](#) model as one where bargaining itself is essentially instantaneous, but payoffs nonetheless are significantly delayed with any disagreement. Based on this interpretation, we generalize the model to arbitrary delays upon disagreement and general time preferences, under the sole substantial assumption of dynamic consistency. This is a special case of bargaining over a time-varying surplus as considered and geometrically analyzed by [Binmore \(1987\)](#), where the variation in surplus derives from non-constant discounting (see also [Coles and Muthoo, 2003](#)). Relative to this prior work, our theoretical contribution consists in showing that, under very mild assumptions on time preferences, there is a unique equilibrium and providing an algebraic proof.

Experiments. [Weg, Rapoport, and Felsenthal \(1990\)](#) and [Rapoport, Weg, and Felsenthal \(1990\)](#) are the first experimental studies of an infinite-horizon, alternating-offers bargaining game. Both implement a within-subjects shrinking-pie design. They compare two conditions, equal and unequal “discount factors,” which correspond to the rates at which the players’ value of the pie shrinks over bargaining rounds.⁹ To prevent their experiments from lasting too long, they program the computer to terminate the bargaining once the number of rounds exceeds 20 while informing their participants only that a game would be terminated by the experimenters if it lasted “too long.” Based on an analysis of their experimental data on final agreements, initial offers, the number of rounds to reach

⁹[Rapoport et al. \(1990\)](#) actually implement fixed costs per round of disagreement rather than constant shrink rates.

agreement and the characteristics of counteroffers, they reject most basic predictions of the [Rubinstein \(1982\)](#) model and argue for the importance of fairness concerns. In particular, they observe neither a significant proposer advantage nor any significant cost advantage.

[Zwick, Rapoport, and Howard \(1992\)](#) experimentally study an environment in which the number of bargaining periods is unlimited and the pie’s value is fixed but bargaining is subject to exogenous random termination. This takes the form of a constant and commonly known breakdown probability. They implement three different such probabilities of breakdown in a between-subjects design. Based on their experimental results, they also reject basic predictions of the [Rubinstein \(1982\)](#) model; e.g., average Round-1 demands are the same under a breakdown probability of 1/10 as under a breakdown probability of 5/6. Furthermore, they reject the equal split solution.

Like [Weg et al. \(1990\)](#) earlier, [Binmore, Swierzbinski, and Tomlinson \(2007\)](#) employ a shrinking-pie design with unequal discount factors. They also adopt a similar forced termination procedure: participants are informed that there will be exogenous termination but not of the exact rule. In fact, the computer intervenes and terminates the game after a randomly drawn number of rounds ranging from 3 to 7. These authors find some behavioral support for the basic predictions of the [Rubinstein \(1982\)](#) model, especially for a proposer advantage. Unlike any of the above studies and ours, however, they have a long and incentivized training/conditioning phase where participants play against a robot programmed to a specific strategy; they also do not implement the deterministic alternating-offers protocol but instead a random proposer protocol, where the proposer of any round is always randomly chosen from the two players with equal probability; moreover, the pie in their experiment consists of lottery tickets.

Notably, none of these studies features any payoff delay. The domain of outcomes over which preferences are defined is either that of immediate monetary rewards or of lotteries over monetary rewards.¹⁰ Hence, none of these studies speaks directly to the question of whether time preferences matter in bargaining; in particular, by their design, they cannot address whether patience is a source of bargaining power, which is the focus of our study, where we implement significant delays to payoffs.¹¹ Another distinctive feature of our study is that we focus on general comparative statics predictions under different theories of time preferences rather than testing particular point predictions against each other.

Regarding our basic question of whether time preferences matter in bargaining, the most closely related work is an unpublished experiment by [Manzini \(2001\)](#). However, her design and also her

¹⁰Somewhat relatedly, [Andreoni and Sprenger \(2012\)](#) conclude from their experimental findings that “risk preferences are not time preferences.”

¹¹Of course, it is straightforward to provide assumptions under which the games thus implemented are formally equivalent to special cases of the [Rubinstein \(1982\)](#) model with exponentially discounted utility. This only means that one may appeal to this equivalence in order to obtain the predictions from that original model under these (more or less stringent) assumptions. It does not mean, however, that one learns anything about how time preferences affect the bargaining outcome; as an analogy, one obviously could not measure time preferences without having any delayed options.

conclusion are radically different from ours. She first elicits participants’ limit prices for avoiding a delay of one and two months, respectively, of a given monetary prize that is otherwise paid the next day, via a variation of the BDM (Becker, DeGroot, and Marschak, 1964) procedure. Then, she pairs the participants for a single bargaining game with alternating offers over just two rounds, so the second round is an ultimatum game. Immediate agreement results in payment the subsequent day, whereas delayed agreement results in payment with a month’s delay.¹² Providing the bargainers with information on their respective limit prices for a month’s delay, this turns out to have no significant correlation with the opening offers.¹³ Hence, she concludes that time preferences do not matter in bargaining and suggests that the task of bargaining distracts attention completely away from time considerations.

Although very carefully designed, the negative conclusion from Manzini (2001) hinges crucially on two assumptions, both related to the participants’ incomplete information regarding their opponent’s preferences. First, it assumes that participants trust the measure of their opponent’s time preference, which includes trust that the opponent understood the elicitation procedure. Second, it assumes that participants’ preferences involve no other concerns confounding time preferences, particularly fairness concerns. We consider these very strong assumptions that are unlikely to hold in the interaction. We also implemented a time preference elicitation task after all bargaining was completed, to check the random assignment in terms of our participants’ underlying time preferences and to see whether conventional measures of time preferences explain bargaining behavior; we find that they hardly do so, in confirmation of the approach we take here (see Appendix F for details). Indeed, we ascribe our very different findings to the way our design deals with incomplete information; by transparently manipulating effective time preferences, our procedure appears very powerful for establishing a setting in which both individuals share an understanding of who is “more patient” and translate this understanding into bargaining power.

3 Theoretical Background

We now present the bargaining game that we implement in our experiment, characterize its unique equilibrium under full generality with regards to time preferences, and show how it is a generalization of the classic Rubinstein (1982) model. In doing so, we highlight two alternative interpretations of the latter in terms of the timing of offers versus payoffs and point out how each of these relates to assumptions about time preferences. All formal proofs are in Appendix A.

¹²Indeed, this is the only other bargaining study we know of that implements delayed payoffs.

¹³She also studies two additional treatments implementing shrinking pies in a way that is comparable to the treatment with delayed payments. For both treatments, she finds much higher correlations of opening offers with the opponent’s cost of disagreeing.

3.1 The Model

Consider two individuals $i \in \{1, 2\}$ deciding on how to share a fixed monetary amount via indefinite alternating-offers bargaining as in [Rubinstein \(1982\)](#). For simplicity, normalize the amount to one, so divisions correspond to shares, and assume it is perfectly divisible. In any round $n \in \mathbb{N}$, one individual i proposes a division $x \in \{(x_1, x_2) : x_1 \in [0, 1] \text{ and } x_2 = 1 - x_1\}$ to the other individual $j = 3 - i$ (we will use this convention for i and j throughout), who can then either accept or reject. If the proposal is accepted, there is agreement, and the game ends; if the proposal is rejected, then the game continues to round $n + 1$, where this protocol is repeated with reversed roles such that j proposes and i responds. Player 1 makes the proposal in round 1, and the game continues until a proposal is accepted. Denoting by r_n the responding player of round n , $r_n = 2$ for n odd, and $r_n = 1$ for n even.

We assume consequentialist individuals who distinguish outcomes only according to whether there is agreement and, if so, how much they obtain at what point in time. Importantly, we decouple rounds and delays, which is the key innovation of our experimental design. While bargaining itself takes essentially no time because offers are so frequent that there is negligible delay between rounds, any disagreement nonetheless entails a significant payoff delay. For the general model, we allow this payoff delay to be arbitrary and to differ between individuals and rounds. Therefore, we specify the domain of individual i 's preferences as any $(q, n) \in ([0, 1] \times \mathbb{N}) \cup \{(0, \infty)\}$, where $(0, \infty)$ subsumes any infinite history (perpetual disagreement). We assume these preferences to have a utility representation U_i that satisfies general discounted utility; i.e., for each individual $i \in \{1, 2\}$, there exist a delay discounting function d_i and an atemporal utility function u_i such that

$$U_i(q, n) = d_i(n - 1) \cdot u_i(q),$$

and that satisfy the following three properties:

1. (Delay Discounting) $d_i(0) = 1 > d_i(n) > d_i(n + 1) > 0 = d_i(\infty)$ for all $n \in \mathbb{N}$;
2. (Atemporal Utility) $u_i : [0, 1] \rightarrow [0, 1]$ is continuous and strictly increasing from $u(0) = 0$ to $u(1) = 1$;¹⁴
3. (Intertemporal Utility) There exists $\alpha_i < 1$ such that for all $n \in \mathbb{N}$, and for all $q \in [0, 1]$ and $q' \in (q, 1]$,

$$u_i^{-1}(\delta_i(n) \cdot u_i(q')) - u_i^{-1}(\delta_i(n) \cdot u_i(q)) \leq \alpha_i \cdot (q' - q),$$

where $\delta_i(n) \equiv d_i(n) / d_i(n - 1)$.

The discounting function $d_i(n - 1)$ gives the discount factor for the total payoff delay associated with agreement being reached in round n , i.e., after $(n - 1)$ rounds of disagreement. The expression $\delta_i(n)$ is the discount factor for the specific period of payoff delay caused by disagreement in round n ;

¹⁴The assumption that $u(1) = 1$ is a mere normalization and without loss of generality.

by property 1, it lies between zero and one. Note that $d_i(n) = \prod_{m=1}^n \delta_i(m)$ holds true, subject to the convention that the “empty product” equals one.

Properties 1 and 2 define the bargaining problem: on the one hand, any round of disagreement causes (further) payoff delay, which is costly to both individuals because they are impatient, and on the other hand, each of them always wants more of the cake for herself.

Property 3 guarantees uniqueness of equilibrium by ensuring that backwards-induction dynamics are well-behaved. It says that i ’s willingness to pay to avoid another round’s delay is always increasing in the amount that she would obtain in case of this delay. This property extends what has been termed “increasing loss to delay” (see the axiomatic formulation of [Rubinstein, 1982](#) and its treatment in [Osborne and Rubinstein, 1990b](#)) or “immediacy” (see the utility formulation of [Schweighofer-Kodritsch, 2018](#)) to the non-stationary setting studied here, and it is implied by standard assumptions; e.g., u_i concave and $\sup_n \delta_i(n) < 1$.¹⁵

To see that the [Rubinstein \(1982\)](#) model is a special case, simply let $\delta_i(n)$ be a constant for each individual i . Given exponential discounting, this means that the payoff delay associated with any round of disagreement is of the same length.¹⁶

3.2 Equilibrium

Our equilibrium notion for this extensive-form game of perfect information is that of subgame perfect Nash equilibrium (SPNE). SPNE outcomes of a more general version of this game, where bargaining is over a general time-varying surplus, are geometrically analyzed by [Binmore \(1987\)](#), who shows that the extreme utilities are obtained in history-independent SPNE. [Coles and Muthoo \(2003\)](#) establish existence for a version of that game, which also contains our model. We contribute here a uniqueness result and a characterization for general discounted utility where non-stationary discounting is the source of time-varying surplus, and we provide algebraic proofs.

Lemma 1. *There exists a unique sequence x_n such that, for all $n \in \mathbb{N}$,*

$$x_n = 1 - u_{r_n}^{-1}(\delta_{r_n}(n) \cdot u_{r_n}(x_{n+1})). \quad (3.1)$$

¹⁵Let u be concave, $q_0 < q_1$ and $\varepsilon > 0$. Then

$$\frac{u(q_0 + \varepsilon) - u(q_0)}{\varepsilon} \geq \frac{u(q_1 + \varepsilon) - u(q_1)}{\varepsilon} > \frac{\delta u(q_1 + \varepsilon) - \delta u(q_1)}{\varepsilon}$$

for any $\delta < 1$. Moreover, if $u(q_0) = \delta u(q_1)$, then $u(q_0 + \varepsilon) > \delta u(q_1 + \varepsilon)$ follows immediately from the above. This is equivalent to $\varepsilon > u^{-1}(\delta u(q_1 + \varepsilon)) - q_0$ and upon substituting $q_0 = u^{-1}(\delta u(q_1))$ to $\varepsilon > u^{-1}(\delta u(q_1 + \varepsilon)) - u^{-1}(\delta u(q_1))$. Denoting $q \equiv q_1$ and $q' \equiv q_1 + \varepsilon$, and applying this to individual i ’s preferences, the third assumed property follows for any given n ; $\sup_n \delta_i(n) < 1$ ensures boundedness away from equality across all n by ruling out that $\lim_{n \rightarrow \infty} \delta(n) \rightarrow 1$.

¹⁶Even when it reduces to that of [Rubinstein \(1982\)](#), the model is not susceptible to the “smallest-units” critique of [van Damme, Selten, and Winter \(1990\)](#) because despite the frequent offers, any disagreement still entails substantial payoff delay.

Proposition 1. *There exists a unique equilibrium. This unique equilibrium is in history-independent strategies that imply immediate agreement in every round. It is characterized by the unique sequence x_n of lemma 1 as follows: in round n , the respective proposer demands share x_n , and the respective respondent accepts a demand q if and only if $q \leq x_n$.*

Proposition 1 delivers a general characterization of SPNE. It has the familiar property that in each round, the proposer makes the smallest acceptable offer to the respondent, given the unique continuation agreement that results from rejection. Hence, in terms of time preferences as of a given round n , only the respondent’s discount factor for that round’s delay $\delta_{r_n}(n)$ enters the equilibrium outcome. In the special case where the model reduces to Rubinstein’s, which will serve as our benchmark, the infinite sequence in (3.1) reduces to two equations:

$$\begin{aligned} x_1 &= 1 - u_2^{-1}(\delta_2 \cdot u_2(x_2)), \\ x_2 &= 1 - u_1^{-1}(\delta_1 \cdot u_1(x_1)). \end{aligned}$$

We generate several behavioral predictions from this exponential-discounting benchmark for our concrete experimental treatments, and we employ the general characterization to also derive the behavioral predictions from various alternative forms of discounting (in particular, quasi-hyperbolic discounting capturing a present bias). We present all of these theoretical predictions in Section 4 after defining our specific treatments.

3.3 Dynamic Consistency and Alternative Interpretation

With frequent offers, the time that passes between decisions is negligible. Hence, preferences are “trivially” dynamically consistent. Essentially, a single self of the individual makes all the strategic decisions, and thus, only this one temporal snapshot of preferences matters (sometimes called “commitment preferences”). This affords our model full generality in terms of time preferences, as it avoids any of the strategic complications arising from dynamic inconsistency that actually manifests itself throughout the game.¹⁷

In fact, dynamic consistency is the only substantial restriction our model imposes: Each individual’s preferences over various outcomes (q, n) are represented by a single utility function U_i as above, and at any point in the game, she consistently maximizes this utility.¹⁸ This means that our model allows

¹⁷For an analysis of dynamically inconsistent but stable time preferences when bargaining itself does take significant time and is therefore subject also to *intra*-personal conflict, see Schweighofer-Kodritsch (2018). However, the general model here can even accommodate dynamically inconsistent preferences that are time-varying because *any* variation over time is negligible for the strategic interaction with frequent offers; e.g., an individual may discount utility exponentially at every point in time but exhibit a higher discount factor on one day than another. For the purposes of our game, her preferences satisfy exponential discounting even though they are dynamically inconsistent across calendar time.

¹⁸We focus on the separable case of discounted utility merely to notationally ease the exposition. It is relatively straightforward to formulate the three assumed properties for non-separable preferences and to then generalize our uniqueness and characterization result using the same line of proof.

not only for arbitrary (though costly) payoff delays upon any disagreement but also delays between rounds. As stressed before, when the latter is negligible, preferences are dynamically consistent by the game’s design. However, under the assumption that preferences are *truly* dynamically consistent, the model accommodates also any setting where bargaining takes a significant amount of time and payoffs may be delayed by the process of bargaining itself. In particular, we may replace rounds n with actual time t , so that there is a significant delay between rounds of bargaining and payoffs occur immediately upon agreement. This is the usual formulation and interpretation of the [Rubinstein \(1982\)](#) model. From this perspective, our model generalizes this classic model from exponential discounting to *any dynamically consistent discounting*.¹⁹

To summarize, the bargaining game we implement in our experiment may appear somewhat artificial, but it has several important practical advantages for our experimental investigation over the usual interpretation of alternating-offers bargaining with a significant delay between rounds of offers (see our introductory discussion in Section 4 below). Moreover, under the assumption of dynamic consistency (or a mere common belief in such consistency), it is also equivalent to a game with the usual interpretation. Proposition 1 and any behavioral predictions derived from it then directly extend to the usual setting where bargaining itself takes time. This applies in particular to the special case of our model with exponential discounting, which is equivalent to [Rubinstein \(1982\)](#) and will therefore serve as our benchmark.

4 Experimental Design and Behavioral Predictions

In this section, we first introduce our general experimental approach to testing predictions of the bargaining theory based on time preferences. We then describe our concrete experimental design and subsequently present the behavioral predictions for the specific treatments. We conclude the section by providing further administrative details.

4.1 General Approach

The theory of bargaining developed here focuses on time preferences as the strategic determinant of bargaining outcomes. When it comes to testing its predictions experimentally, a researcher first faces the challenge that she does not know the participants’ preferences. Equation (3.1) demonstrates that the theory’s point prediction is potentially influenced by an infinite number of discounting parameters together with the shape of the atemporal utility functions, whose elicitation would not be practically

¹⁹Non-exponential but dynamically consistent discounting means that discounting may vary with absolute (calendar) time but not the passage of time; e.g., an individual’s discount factor for the first week of January 2022 may differ from her discount factor for the second week of January 2022, but both of these discount factors are the same regardless of when the individual considers these delays. See [Halevy \(2015\)](#) for further elaboration and experimental investigation of the relationships between stationarity, dynamic consistency and time invariance (and their respective violations).

feasible. Moreover, preferences most likely include concerns other than time preferences, particularly fairness concerns (as demonstrated by the ultimatum game; see [Güth, Schmittberger, and Schwarze, 1982](#)). For this reason, we design our experiment to test comparative statics predictions regarding time preferences rather than point predictions. In other words, we ask whether time preferences matter as predicted despite various other possible concerns.

Second, and in contrast to the theory, the players themselves also do not know their opponent’s (time, risk, social) preferences. Put simply, the basic goal of our design is to establish a setting where both individuals have a shared understanding of who is more “patient,” which we would argue is the setting to which the theory is really meant to apply. Our key innovation towards this goal is to randomly assign participants their individual payoff delay structures and to make this assignment commonly known within each matched bargaining pair. Thus, we create groups of individuals with the same distribution of *general* preferences (including various concerns) but effectively different *time* preferences. Assumptions on the structure of *underlying* time preferences then translate into shifts in relative bargaining power via predicted shifts in patience, and making both individuals’ “types” common knowledge then implies common beliefs about who will be at an advantage in terms of their *effective* time preferences. We derive such comparative statics predictions for the most important classes of time preferences without any parametric assumptions, and we test them by comparing the associated distributions of behavioral outcome measures for different types.

4.2 Experimental Design

Table 1 presents our experimental design, which contains three experimental treatments. Each treatment corresponds to a particular pairing of “bargainer types,” where this type corresponds to the exogenously imposed payoff delay that an individual faces for any possible disagreement. To credibly implement meaningful payoff delay, we relied on the popular mobile payment system *Venmo*.²⁰

In Treatment *WM*, one bargainer faces one week of delay per round of disagreement, whereas the other faces one month of such delay. Treatment *WM2D* is similar, but both bargainers additionally face a front-end delay of one week; i.e., now immediate agreements also result in one week of payoff delay. In Treatment *WW1D*, both bargainers face the same delay of one week per round of disagreement, but one of them additionally faces a front-end delay of one week. In the rest of the paper, we will call the bargainer whose payment window is weekly/monthly/delayed a *weekly/monthly/delayed* bargainer.

Within a given treatment, all games are played by a particular pair of different types, and everyone

²⁰Venmo is a service provided by PayPal that allows account holders to transfer funds to others via a mobile phone app. It handled \$12 billion in transactions during the first quarter of 2018 (<https://en.wikipedia.org/wiki/venmo>). For more information, please visit <https://help.venmo.com/hc/en-us/articles/210413477>. When recruiting our participants, we clearly announced that those without a Venmo account were not eligible to participate in the experiment. At the end of the experiment, the participants were asked to report their account information for payment, including username and email address details. None of the participants reported any error or difficulty in providing this information, suggesting that all our participants were sufficiently familiar with Venmo in their daily lives.

Table 1: Experimental Treatments

Bargainer 1	Bargainer 2		
	Monthly with D	Monthly	Weekly with D
Weekly with D	WM2D	N/A	N/A
Weekly	N/A	WM	WW1D

*Note: Delay (D) = 1 Week

anonymously plays ten games. Participants are always randomly rematched subject to the treatment condition. Moreover, the initial proposer is always determined by chance, so we observe both versions of the game in terms of which type is the initial proposer. For instance, in every game of Treatment *WM*, a weekly bargainer plays against a monthly bargainer, and half of the games have a weekly bargainer as the initial proposer and a monthly bargainer as the initial respondent; the other half have a monthly bargainer as the initial proposer and a weekly bargainer as the initial respondent. We can compare these two kinds of games within treatments to measure and test for a basic proposer advantage.

Our focus is on testing comparative statics with respect to time preferences, however, and we now sketch how our treatments deliver such tests. Details and formal derivations of the behavioral predictions to be tested follow below in Section 4.3. Whenever a weekly and a monthly bargainer are matched, one transparently faces a longer delay and therefore greater cost of disagreement than the other for all commonly considered time preferences. Thus, we can test whether effectively greater patience translates into a strategic advantage in terms of a more favorable bargaining outcome. Introducing a front-end delay additionally allows us to test the prediction from exponential discounting that only a “marginal” delay but not a “fixed” delay matters—akin to marginal v. fixed/sunk cost—against the alternative of present bias or also future bias. Observe here that our treatments produce such tests not only within treatments (across the two kinds of games) but also across treatments (weekly bargainers appear in both *WM* and *WW1D*, and delayed weekly bargainers appear in both *WM2D* and *WW1D*).

In the following, we highlight the key components of our experimental design and compare them with the conventional designs used in the related literature. The full experimental instructions for Treatment *WM* can be found in Appendix D.

Effective Discounting Procedure and Payoff Delay. Unlike the shrinking-pie design, the size of the surplus in our experiment is fixed, at US\$50 (500 tokens), and we use the novel experimental manipulation proposed by Kim (2019b), the “effective discounting procedure,” to control time preferences and implement payoffs over a potentially long period of time.²¹ More precisely, we exogenously control

²¹Kim (2019b) develops the design to investigate the effect of time preferences on cooperative behavior in an infinitely repeated game.

the *effective* discounting of our participants by changing the payment delay of bargaining payoffs at the individual level, which is either a week or month per round to agreement and may include an additional front-end delay of one week. For instance, in Treatment *WM*, the weekly bargainer receives the payoff from an agreement in Round n in $(n - 1)$ week(s) from the day of the experiment; the monthly bargainer receives her payoff in $(n - 1)$ month(s) from the day of the experiment. Appendix D illustrates how the payment schedule is presented to participants in this treatment. With the additional front-end delay in Treatment *WM2D*, the now delayed weekly and monthly bargainers receive their payoff in n week(s) and in $(n - 1)$ month(s) plus one week, respectively. Assuming (β, δ) -discounting (i.e., quasi-hyperbolic discounting), any present (or near-future) bias will vanish with such a front-end delay, and this will also be true with more general forms of such a bias.²² Our choice of basic delays of a week versus a month intends to ensure salient differences in bargainers’ effective discounting in the respective treatments, in spite of interests rates that are close to zero and of the aforementioned likely relevance of concerns other than time preferences.

Fixed Types and Random Rematching. To minimize any potential confusion among participants regarding their incentives and payment schedule, we randomly assigned every participant their “bargainer type” at the beginning of the session and fixed it throughout the entire experiment. The participants then played multiple games as usual, where they were randomly rematched after each game, subject to the particular treatment’s pairing specification. Within any match, the types of both players were made common knowledge at the very beginning of the game; see the screenshots in Appendix E.

Probabilistic Termination. In addition to payoff delay, we also implemented exogenous termination with a fixed, commonly known termination probability of 25%. By contrast, previous studies of the Rubinstein (1982) model, such as Zwick et al. (1992), have employed probabilistic termination as the sole cost of disagreement and studied the effects of varying the termination probability on outcomes.²³ Importantly, in our experiment, the same 25% termination probability was transparently applied to all rounds of all games in all treatments. As a result, even if risk attitudes enter preferences, they could not be a significant source of the behavioral differences we observe. This design choice serves two related purposes. First, it ensures that every bargaining game, while still indefinite, is expected to end after a reasonable amount of time and, upon its conclusion, to be promptly followed by the next one, which is important for the credibility as well as the smooth running of our experiment. Second, this design theoretically keeps bargainers further away from possible indifference to delay as required by Property 3 of our preference assumptions. Of course, in terms of our model, discounting should be viewed as also including this constant risk (assuming expected utility).²⁴

²²There is no consensus on how long this delay needs to be to make any present or (near-) future bias negligible. In his related experiment on the repeated Prisoners’ Dilemma, Kim (2019b) finds that the cooperation rate is significantly higher in a treatment with a payoff delay of one month per round than in a treatment with no payoff delay, suggesting that a month may be an upper bound.

²³See also Dal Bó and Fréchette (2018) for the use of the probabilistic termination in infinitely repeated games.

²⁴With expected utility, a constant probability of breakdown simply proportionally reduces each $\delta_i(n)$ by this fraction,

Deterministic Roles. At the beginning of the first round of every bargaining game, the initial proposer/responder roles were randomly determined, which in turn determined the individuals’ roles for all subsequent rounds via the alternating-offers protocol. This is in contrast to experiments studying random-proposer bargaining protocols, where the proposer is randomly determined in *every round*; [Binmore et al. \(2007\)](#) is the most closely related such experiment.

4.3 Behavioral Predictions

We now employ Proposition 1 to derive the behavioral predictions that our experiment is designed to test. These concern the basic theoretical implications of the most important classes of time preferences with regards to efficiency and distribution. Efficiency of bargaining (immediate agreement) is a general implication. The distribution of surplus (bargaining power) depends on who is the initial proposer (proposer advantage) and on the two bargainers’ specific time preferences (“patience advantage”). All formal proofs are in Appendix A.

We begin by establishing the important and influential benchmark predictions from exponential discounting, as in [Rubinstein \(1982\)](#), and subsequently highlight the differential predictions under the most important alternative forms of discounting as observed empirically – in particular, present bias as in quasi-hyperbolic discounting. In each case, to capture the implied “typical” behavior, we impose preference symmetry: i.e., both individuals have the same atemporal utility function, $u_1 = u_2 = u$, and for the same future delay $\Delta_{t,t'}$ from some given date $t > 0$ to some later date $t' > t$, discount utility with the same discount factor $\delta_{t,t'}$. Note that by implementing idiosyncratic payoff delays (bargainer types), our effective discounting procedure nonetheless induces variation in the cost of disagreement within and across matches/treatments.

Exponential Discounting (EXD). Since any given bargainer type faces the same payoff delay from any round of disagreement, the stationarity property of EXD implies that any such delay is discounted with the same discount factor, irrespective of any front-end delay. Let $\delta \in (0, 1)$ be the (common) discount factor for a weekly delay, and let $\phi\delta$ be the (common) discount factor for a monthly delay, where $0 < \phi < 1$.²⁵ Using notation $\phi_i \in \{\phi, 1\}$ with $\phi_i = 1$ if and only if bargainer i is a weekly bargainer, any bargainer i ’s type is fully captured by ϕ_i , such that $U_i(q, n) = (\phi_i\delta)^{n-1}u(q)$ and $\delta_i(n) = \phi_i\delta$ is constant across rounds n . Both *WM* and *WM2D* correspond to pairing $\{1, \phi\}$, and *WW1D* corresponds to pairing $\{1, 1\}$.

whereby it is captured by our model. It should be noted, however, that certain violations of expected utility yield dynamic inconsistency even across *rounds*, in contrast to our model. In particular, [Halevy \(2008\)](#) argues that the future is inherently uncertain and shows how empirically plausible non-linear probability weighting of future consumption risk provides a foundation of present bias and diminishing impatience. Since [Schweighofer-Kodritsch \(2018\)](#) finds that this form of dynamic inconsistency does not upset any of the qualitative equilibrium predictions of the benchmark under dynamic consistency, we abstract from risk as a potential source of dynamic inconsistency.

²⁵If we take a month to equal four weeks, then $\phi\delta = \delta^4$ pins down $\phi = \delta^3$.

Prediction 1. Symmetric EXD implies:

- (1) **Efficiency:** There is always immediate agreement.
- (2) **Proposer Advantage:** In every treatment, a given bargainer type obtains a greater share as the initial proposer than as the initial respondent.
- (3) **Basic Delay Advantage:** For a given initial role (proposer or respondent):
 - (3a) In both Treatments *WM* and *WM2D*, the weekly bargainer obtains a greater share than the monthly bargainer.
 - (3b) Across Treatments *WM* and *WW1D*, the weekly bargainer obtains a greater share against the monthly bargainer (*WM*) than against the delayed weekly bargainer (*WW1D*).
 - (3c) Across Treatments *WM2D* and *WW1D*, the delayed weekly bargainer obtains a greater share against the delayed monthly bargainer (*WM2D*) than against the non-delayed weekly bargainer (*WW1D*).
- (4) **Front-End Delay Neutrality:** For a given initial role (proposer or respondent):
 - (4a) In Treatment *WW1D*, both types of bargainer obtain the same share.
 - (4b) Across Treatments *WM* and *WM2D*, bargainer types with the same basic delay (weekly or monthly) obtain the same share.

Efficiency (1) and Proposer Advantage (2) are well-understood predictions. For the comparative statics predictions in effective time preferences, (3) and (4), simply note that with EXD front-end delay is strategically irrelevant, so that the Treatments *WM* and *WM2D* become equivalent and Treatment *WW1D* becomes symmetric; the predictions are then an immediate consequence of the fact that any weekly bargainer is effectively more patient than any monthly one.²⁶

Quasi-Hyperbolic Discounting (QHD). Present bias, the excessive weight put on immediate rewards relative to delayed rewards, is the most important deviation from EXD. By adding a single parameter $\beta \in (0, 1)$, the model of quasi-hyperbolic discounting parsimoniously captures this empirically well-established phenomenon. Given that all offers are made at the same date, the bias may play a role only in the first round, because upon failure to agree immediately all possible payoffs lie in the future. Moreover, it will do so only when the initial respondent faces no front-end delay, because the proposer's discounting of the first round's delay is irrelevant anyways, due to the proposer's strategic advantage, and a front-end delay for the respondent pushes any immediate-agreement payoff into the

²⁶It seems worthwhile pointing out that whereas the Proposer Advantage (2) concerns a comparison across the two possible initial role assignments for the same type, the comparative statics predictions (3) and (4) concern comparisons across types for the same initial role assignment.

future. Keeping the earlier EXD notation and adding $\beta_i \in \{\beta, 1\}$ with $\beta_i = 1$ if and only if bargainer i is not delayed, any bargainer i 's type is fully captured by (ϕ_i, β_i) , such that $U_i(q, n) = \beta_i (\phi_i \delta)^{n-1} u(q)$; now $\delta_i(1) = \beta_i \phi_i \delta$ and, for $n > 1$, $\delta_i(n) = \phi_i \delta \geq \delta_i(1)$. *WM2D* corresponds to pairing $\{(1, 1), (\phi, 1)\}$, *WM* corresponds to pairing $\{(1, \beta), (\phi, \beta)\}$, and *WW1D* corresponds to pairing $\{(1, \beta), (1, 1)\}$.

Prediction 2. Symmetric QHD implies (1) and (2) as under EXD, and

(3') **Qualified Basic Delay Advantage:** (3a, 3b) as under EXD, and

(3c') Across Treatments *WM2D* and *WW1D*, the delayed weekly bargainer obtains a greater share against the delayed monthly bargainer (*WM2D*) than against the non-delayed weekly bargainer (*WW1D*) as the initial respondent, but *the respective ranking for the delayed weekly bargainer's share as the initial proposer is ambiguous in general*.

(4') **Front-End Delay Advantage:**

(4a') In Treatment *WW1D*, the *delayed weekly bargainer obtains a greater share than the weekly bargainer*, both as initial proposer and as initial respondent.

(4b') Across Treatments *WM* and *WM2D*, bargainer types with the same basic delay (weekly or monthly) obtain a *greater (resp., smaller) share in WM than in WM2D as initial proposer (resp., respondent)*.

Immediate agreement, hence Efficiency (1), is a general prediction. For the remaining ones, note that from Round 2 (which is off-path, of course) the game proceeds as under EXD. Hence, a present bias in the sense of $\beta < 1$ kicks in only if the initial respondent is not delayed, in which case it could only reinforce the Proposer Advantage (2). Due to this advantage, the initial proposer's β does not affect the equilibrium agreement, whereby any front-end delay to the initial proposer's payoff has no effect under QHD.

For (3a), the prediction within Treatment *WM2D* is therefore immediate from EXD. Within Treatment *WM*, a weekly bargainer as initial respondent would obtain a better deal after rejecting than a monthly one would and with the same present bias is more patient about the first delay; therefore, she obtains a greater share in the immediate agreement, where the proposer's offer always makes the respondent indifferent to rejecting. A similar argument applies when we compare a weekly bargainer's share as initial respondent against a monthly bargainer (*WM*) with that against a delayed weekly bargainer (*WW1D*), because rejection results in a better deal in the former case; for the remaining part of (3b) and the weekly bargainer as initial proposer, note that a present bias makes a monthly respondent even weaker relative to the delayed weekly respondent.

The part of (3c') that coincides with EXD's (3c) has the initial respondent delayed, so it is indeed immediate from (3c). Now compare a delayed weekly bargainer as the initial proposer against a delayed monthly bargainer (*WM2D*) and against a weekly bargainer (*WW1D*). Although the Round-2

(off-path) continuation agreement would have the monthly one weaker due to her longer basic delay, present bias only affects the (non-delayed) weekly respondent and may overturn the EXD prediction. The comparison therefore depends on how strong present bias β is relative to long-run discounting δ .

Finally, both parts of Prediction (4') are straightforward from the fact that a delayed bargainer is stronger than a non-delayed bargainer with the same basic delay. The symmetry in (4a) within Treatment *WW1D* under EXD is therefore broken in favor of the delayed weekly bargainer (4a'). The logic of part (4b') across Treatments *WM* and *WM2D* follows similarly, upon recalling that the initial proposer's present bias β is irrelevant.

Other Forms of Discounting. Due to the tractability they afford, EXD and QHD are, by far, the most important models of time preferences for theoretical analyses. However, empirical studies, especially from psychology, suggest hyperbolic discounting—also known as diminishing impatience, which implies present bias—as more “universal” form of discounting. At the same time, experimental studies from economics also document the “opposite” of present bias, namely (near-) future bias. We now discuss the implications of these alternatives.

First, consider hyperbolic discounting, where $\delta_i(n)$ is increasing in n . Since it implies a present bias, a front-end delay increases such a discounter's bargaining power as the respondent. However, disagreement in round n adds a shorter payoff delay to a shorter delay for a weekly bargainer than a monthly bargainer, meaning that for n large enough a monthly bargainer may in general become more patient than a weekly bargainer. This would resonate through the entire recursion of equation (3.1), thereby affecting the equilibrium outcome. Based on intuition that discounting for the same additional delay would not change too quickly with the preceding delay (except for the immediate present) and the sizable termination probability, we assume that the effect of pushing delays further into the future does not outweigh the effect of longer delays in determining the immediate equilibrium agreement. Notably, the leading models of hyperbolic discounting are all special cases of the discounting function $d(t) = (1 - \alpha \cdot t)^{-\beta/\alpha}$ (with $\alpha, \beta > 0$) proposed by [Loewenstein and Prelec \(1992\)](#), and for any such discounting our design ensures that the cost of disagreement is always lower for a weekly bargainer than for a monthly bargainer (also when *both* are delayed).²⁷ Loosely speaking, hyperbolic discounting then remains sufficiently similar to constant discounting for any delay that starts in the future, such that by continuity, the only qualitative difference from EXD *within* our treatments is a present bias, as captured by QHD. Across treatments, however, hyperbolic discounting still introduces another possibility: even when restricted to the above functional form, it remains ambiguous whether the weekly (with no delay) or the delayed monthly bargainer type is more patient (although one can show

²⁷The reason is that the different delays per round have a constant ratio, which also equals the ratio of total delays the two bargainers face in any agreement. Measuring time t in the unit that is the shorter delay per round and letting the corresponding type be type A , A 's discount factor for round n is $\delta_A(n) = [(1 - \alpha \cdot n)/(1 - \alpha \cdot (n - 1))]^{-\beta/\alpha}$; letting the longer delay be $k > 1$ times the shorter delay with corresponding type B , B 's discount factor for round n is $\delta_B(n) = [(1 - \alpha \cdot kn)/(1 - \alpha \cdot k(n - 1))]^{-\beta/\alpha}$. Basic algebra yields $\delta_A(n) > \delta_B(n)$, and it is straightforward to check that the same holds true if both A and B face the same front-end delay.

that they can always be ordered). This renders hyperbolic discounting altogether permissive with respect to Predictions (3c/3c') and (4b/4b').

Finally, consider also near-future bias. Somewhat loosely, this means that the discounting function is initially concave (hump-shaped), in contrast to the convex discounting functions under EXD, QHD or hyperbolic discounting. While empirically documented, it is neither known how prevalent this bias is (hence, whether it could be reasonably expected to guide typical behavior) nor how far the “near” future extends from the immediate present (hence, whether a week’s front-end delay could be reasonably expected to mute it). In view of these open issues, we omit a detailed analysis but note that if a near-future bias operates like “inverted” present bias in QHD—i.e., $1 < \beta < 1/\delta$ —then it would simply yield the mirror image of the differential predictions of QHD from EXD because a front-end delay would then make the initial respondent weaker rather than stronger.²⁸

Summary Immediate agreement and a proposer advantage are fundamental theoretical implications for bargaining. While our experiment is set up to test these as well, its main innovation and focus concern comparative statics in time preferences. The idea that patience is power in bargaining becomes somewhat complex with any violation of EXD since there is then no simple measure of patience to apply at every stage of bargaining. As argued, under mild restrictions, a basic delay advantage in the sense of prediction (3a) may be considered a fundamental theoretical implication as well. Focusing on QHD as parsimoniously capturing the key empirical deviation from EXD (among all forms of present-biased discounting), the major distinctive implication for our experiment to test is that a front-end delay increases an initial respondent’s bargaining power (rather than having no effect), which underlies all of its differential predictions. Under minimal restrictions, hyperbolic discounting maintains the fundamental predictions together with the key distinctive prediction of QHD but is otherwise more permissive.

4.4 Administrative Details

Our experiment was conducted using z-Tree (Fischbacher, 2007) at the University of California, Irvine. A total of 348 subjects who had no prior experience with our experiment were recruited from the graduate and undergraduate student population of the university. Upon arrival at the laboratory, the participants were instructed to sit at separate computer terminals. Each received a copy of the experiment’s instructions. To ensure that the information contained in the instructions was induced as public knowledge, these instructions were read aloud, and the reading was accompanied by slide illustrations followed by a comprehension quiz.

Each session employed a single treatment, and we conducted 6 sessions for each treatment, for a

²⁸Specifically, rather than qualifying prediction (3c), which would then carry over from EXD, such bias would qualify (3b) to not necessarily hold anymore for the weekly bargainer as the initial proposer, and it would change “greater” to “smaller” in predictions (4a') and (4b').

total of 18 sessions (6 sessions \times 3 treatments). In all sessions, the participants played 10 games under the corresponding treatment condition, say matching bargainer types A and B. At the beginning of the experiment, one half of the participants were randomly assigned to be Type A and the other half to be Type B. Individual participants' types remained fixed throughout the session. We used the random-matching protocol (across matches, subject to the treatment condition). Each session had 16–20 participants and hence involved 8–10 simultaneous games.

We illustrate the instructions with those for Treatment *WM*. The full instructions for this treatment can be found in Appendix D. For each game, one Type A participant and one Type B participant were randomly matched. At the very beginning of Round 1, one of the two was randomly chosen to be the proposer, and the other was the responder. The proposer then proposed how to split 500 tokens (worth \$50) between the two of them as follows:

“----- tokens for yourself and ----- tokens for the other person.”

After observing the proposed split, the respondent decided whether to accept or reject it. If the respondent accepted the proposed split, both participants received their proposed amount in tokens, and the match was terminated. If the respondent rejected the proposed split, then the match proceeded to the next round of bargaining with a 75% chance and was terminated with a 25% chance. If a match was terminated after the rejection of a proposed split, both participants received zero tokens for the match. If the match proceeded to the next round, then the participant who was the proposer in the previous round became the respondent, and the participant who was the respondent in the previous round became the proposer. At the end of the experiment, one of the 10 matches was randomly selected for payment. For the selected match, if agreement was reached, the delay of the participant's payment depended on (1) his/her bargainer type and (2) the round of the agreement.

After all ten matches were over, we measured the participants' time preferences by using a version of the BDM (Becker, DeGroot, and Marschak, 1964) method. We elicited switching points (indifferences) between sooner and later money amounts. One decision was randomly selected for actual payment.²⁹

The tokens our participants earned in the selected match were converted into US dollars at a fixed and known exchange rate of \$0.1 per token. In addition, participants received a show-up payment of \$10. Any amount a participant was due to receive was paid electronically via Venmo, including immediate payments. Earnings were \$37.90 on average, and average duration of a session was approximately 1.5 hours.³⁰

²⁹We implemented the elicitation task in 4 sessions per treatment. This allows us to check whether the random assignment was successfully implemented in terms of participants' underlying time preferences, which is a crucial aspect of our design and which our data confirm. In line with Manzini (2001), we find no significant correlations of these measures with behavior in our experiment, supporting our approach. See Appendix F for details.

³⁰We conducted 6 sessions in May and June, 2018, and 12 sessions in October and December, 2018. The longest delay among the matches selected for payment was 7 months, and the corresponding amount was paid on May 17, 2019.

5 Experimental Results

This section presents our experimental results regarding Predictions 1 and 2. First, we take a look at the basic efficiency property of bargaining outcomes and present empirical evidence regarding the proposer advantage. Then, we investigate the key predictions regarding any advantage created from the basic delay and front-end delay manipulation. In the body of this paper, we mainly focus on the data from the second half of the experiment, after some learning has taken place. In line also with the literature, we conduct our tests based on comparisons of initial proposals to capture the bargainers’ perceptions of relative bargaining power. We have these data for every bargainer type in every treatment, and every participant with a given type played both versions of the game against the treatment’s given opponent type (i.e., as the initial proposer and as the initial respondent). Notably, the vast majority of agreements were reached immediately or with only one round of delay, with highly similar rates across all treatments. In any case, we do find similar results for the first half of the experiment and when comparing actual immediate agreements; Appendix B includes the supporting supplementary figures. Finally, we summarize and discuss our findings.

5.1 Efficiency and Proposer Advantage




Figure 1: The Proportions of Agreements over Rounds – Last 5 Matches

Figure 1 depicts the proportions of agreements made in Rounds 1, 2, 3, and 4 or after based on the data from the last 5 matches. The basic picture does not change if we instead use the first 5 matches only or all 10 matches. In Treatment *WM*, the vast majority of matches result in agreement reached

with no delay (72.5%) or one round of delay (12.5%). In Treatments *WM2D* and *WW1D*, the broad picture is highly similar. For all rounds, the proportions of agreement before random termination are 91.8%, 89.3% and 90.5% for Treatments *WM*, *WM2D*, and *WW1D*, respectively.

In all pairwise comparisons between treatments, the acceptance rates in Round 1—i.e., the rates of immediate agreement—do not significantly differ between treatments (one-sided Fisher’s exact test, p -values > 0.18).³¹ Figure 10 in Appendix B shows that these rates also do not differ significantly across the two versions of the game within any treatment.

These observations establish that immediate agreement is the rule rather than the exception, and they suggest that failures to agree are unrelated to the specific pairing of bargainer types and effective time preferences. Of course, participants make their decisions based on incomplete information, but this does not seem to cause much delay and inefficiency. The fact that the average number of rounds for agreement is only slightly above 1.3 also supports that, on average, bargaining achieved very high levels of efficiency.³²

Finally, considering agreements only (recall the risk of exogenous termination) and averaging also over all treatments, more than 75% are immediate agreements, and approximately 17% are agreements with a delay of only one round. We summarize this as follows.

Result 1 (Efficiency). *In every treatment, the vast majority of agreements are reached immediately or with only one round of delay.*

We next explore the proposer advantage in our data based on initial proposals. Figure 2’s left panel reports the average share for proposers and respondents of each type in each treatment over the last 5 matches. For every type, the average share for proposers is significantly larger than that for respondents (Kolmogorov-Smirnov test, p -values < 0.01).³³ Moreover, the differences are substantial in their magnitudes (25–40 tokens).

Of course, not all proposals are accepted, and we would naturally expect more rejections for proposals that leave less to respondents. Hence, we also consider actually accepted proposals, see figure 2’s right panel for the last 5 matches, which confirm the advantage of being the initial proposer.³⁴ These observations firmly support the predicted proposer advantage in bargaining when the cost of disagreement is delay.

Result 2 (Proposer Advantage). *In every treatment, given any bargainer type (weekly, monthly, delayed), the average share is significantly and substantially larger as proposer than as respondent.*

³¹For a robustness check to control for the size of proposals in Round 1, we also run probit regressions in which a treatment dummy variable and the respondents’ share are independent variables and the standard errors are clustered at the session level. For all pairwise comparisons over the last 5 matches, the difference in acceptance rates is insignificant except for being only marginally significant in the comparison of Treatments *WM* and *WW1D* (p -value = 0.059). For all matches, no pairwise comparison results in a significant difference (p -values > 0.153).

³²The average number of rounds for agreement does not differ across treatments (Mann-Whitney test, p -values > 0.5).

³³Qualitatively the same patterns are observed for the first 5 matches.

³⁴We find a similar confirmation for the first 5 matches and also for actual payoffs; see appendix B.

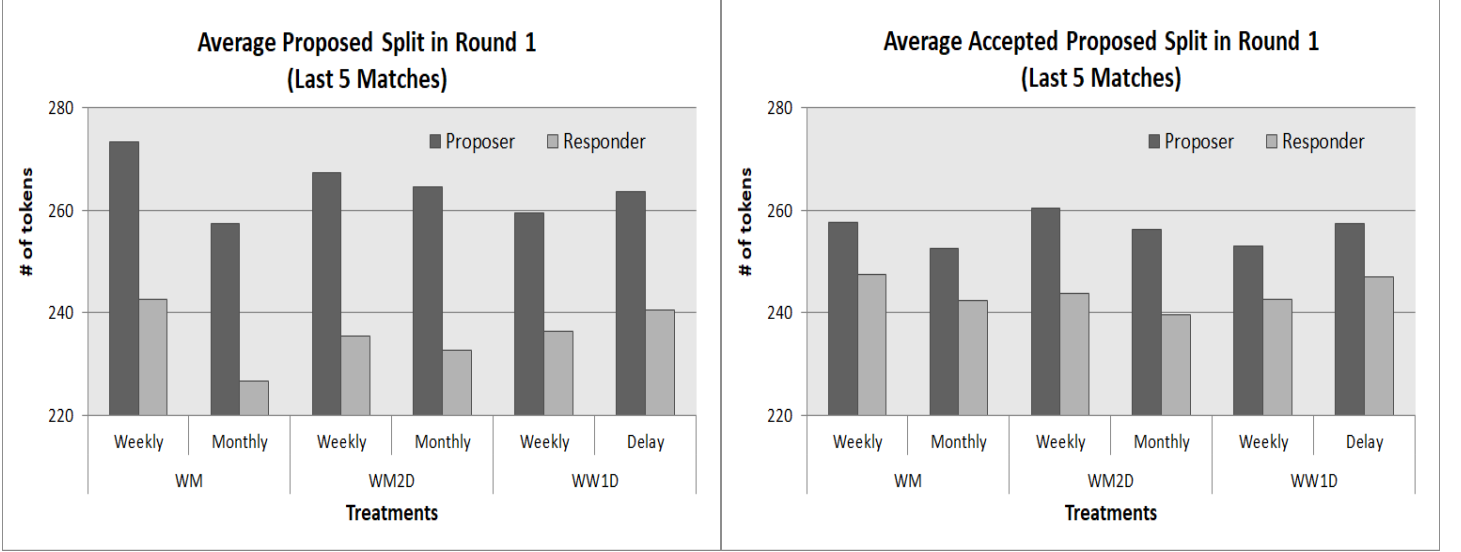


Figure 2: Proposer Advantage – Last 5 Matches’ Proposals

5.2 Basic Delay and Front-End Delay Advantages

To maximize the amount of data for our analyses, we focus on initial proposals, thereby treating them as equilibrium proposals. Contrary to the equilibrium assuming perfect information, not all proposals are accepted, of course, due to preference heterogeneity and incomplete information. However, most are and at very similar rates across conditions. This suggests that the proposers do generally attempt to reach agreement immediately and that the effect of incomplete information is “constant” across conditions. Moreover, we obtain similar results when considering actually agreed shares; see Figures 17 through 19 in Appendix B.

Our manipulation is supposed to effectively shift the distribution of time preferences between the randomly selected groups of bargainer types. We therefore conduct our comparisons based on the entire observed distributions of initial proposals in terms of the proposer’s claimed share.³⁵ Specifically, we always examine the cumulative distribution functions (CDFs) for first-order stochastic dominance, and we use appropriate unidirectional Kolmogorov-Smirnov (KS) tests for statistical significance in this strong sense. Note that in contrast to comparisons of means, which are always ordered, such order is not guaranteed here.

Recall that our design produces tests of all predicted comparative statics on within-treatments data, which we consider first, and on across-treatments data, which follows second.

³⁵The CDF figures that follow below are censored by the given range of [250, 310] for ease of graphical representation. This range contains, on average, more than 95% of the observations in the data. Such truncation does not apply to any of the statistical analyses.

5.2.1 Within Treatments

Within treatments, we test for the basic delay advantage according to Prediction (3a), which is common to both EXD and QHD and concerned with Treatments *WM* and *WM2D*: in each of these treatments, proposals by weekly bargainers are predicted to exceed those by their monthly counterparts in terms of the proposer’s own share. The other within-treatment test is that of front-end delay neutrality under EXD, Prediction (4a), vs. front-end delay advantage under QHD, Prediction (4a’). This test is concerned with treatment *WW1D*: whereas both bargainer types are predicted to make identical proposals under EXD, under QHD the delayed type is predicted to claim a larger share of the surplus.

Note that if a Type A’s initial proposals have a larger proposer share than a Type B’s proposals, this is equivalent to Type A’s being offered a larger respondent share than Type B. Hence, within each treatment, a single comparison covers the entire respective prediction.

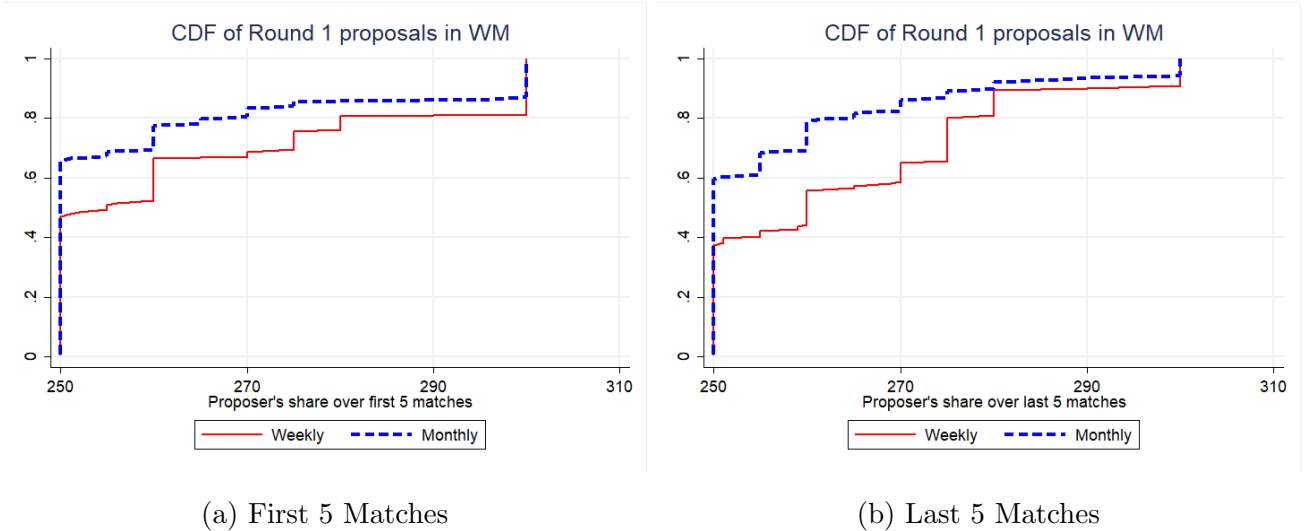


Figure 3: Round-1 Proposals in Treatment *WM*

Figure 3 presents the CDF of Round-1 proposals in Treatment *WM* aggregated over the first 5 matches (Figure 3(a)) and over the last 5 matches (Figure 3(b)), by bargainer type. The solid line indicates the CDF for the weekly proposer, and the dotted line indicates the CDF for the monthly proposer. A few observations are immediate. First, consistent with prior findings, fairness concerns seem to be an important factor in bargaining.³⁶ Approximately 50% of proposals are equal (250-250) splits. Second, the CDF of proposals by weekly bargainers clearly lies below that for monthly bargainers already for the first 5 matches (KS test, p -value = 0.046). This difference remains statistically significant and becomes even more substantial in magnitude in the last 5 matches (KS test, p -value < 0.01). We therefore strongly confirm the general Prediction (3a) for Treatment *WM*. Both matched types act in accordance with a shared understanding that a longer delay increases the cost of disagreement and

³⁶In fact, this does not require the proposers to be fair-minded; it could be that they have “selfish” preferences but believe that they are facing a fair-minded respondent.

weakens bargaining power.

Result 3 (Basic Delay Advantage in Treatment *WM*). *In Treatment WM, the initial proposals by weekly bargainers significantly exceed and first-order stochastically dominate those by monthly bargainers.*




Figure 4: Round-1 Proposals in Treatment *WM2D*

Figure 4 presents the CDF of Round-1 proposals in Treatment *WM2D* by bargainer type. The solid line indicates the CDF for the (now delayed) weekly proposer, and the dotted line indicates the CDF for the (now delayed) monthly proposer. Again, close to 50% of proposals are equal splits. Unlike in Treatment *WM*, the distributions of proposals are quite obviously not significantly different initially (KS test, p -value = 0.726). However, behavior gravitates towards the theoretical prediction as the participants gain more experience. In the comparison for the last 5 matches, we observe the predicted first-order stochastic dominance, although it remains statistically non-significant (KS test, p -value = 0.262). Restricting attention to proposals that are strictly greater than 250 (meaning that the proposer claims more than half the surplus), the predicted first-order stochastic dominance relationship becomes marginally significant (KS test, p -value=0.078).³⁷ Taken together, these findings also support the general Prediction (3a), although relative to Treatment *WM*, the front-end delay in Treatment *WM2D* significantly mitigates the basic delay advantage. In other words, it is less clear that the two bargainer types perceive the longer delay as significantly weakening bargaining power. This observation is consistent with diminishing impatience, so adding a front-end delay shifts bargaining power somewhat to the bargainer with longer delays.

³⁷Excluding the observations with equal shares does not harm our analysis for Treatment *WM2D* because the frequencies of equal-split proposals by weekly and monthly bargainers are not statistically different (Fisher's exact test (one-sided), p -value = 0.378).

Result 4 (Basic Delay Advantage in Treatment *WM2D*). *In Treatment WM2D, the initial proposals by delayed weekly bargainers and those by delayed monthly bargainers are not significantly different in their distributions. We observe first-order stochastic dominance of the former in the last 5 matches, where the difference is statistically more significant, especially upon excluding equal splits.*

While Treatments *WM* and *WM2D* pair two types that differ solely in their payoff delay per round of disagreement and therefore permit a straightforward test of the basic delay advantage, Treatment *WW1D* is symmetric in this respect. The only asymmetry between types here is that one is facing a front-end delay while the other is not. Under EXD, this “fixed cost” asymmetry is irrelevant, while it is an advantage under QHD (present bias).

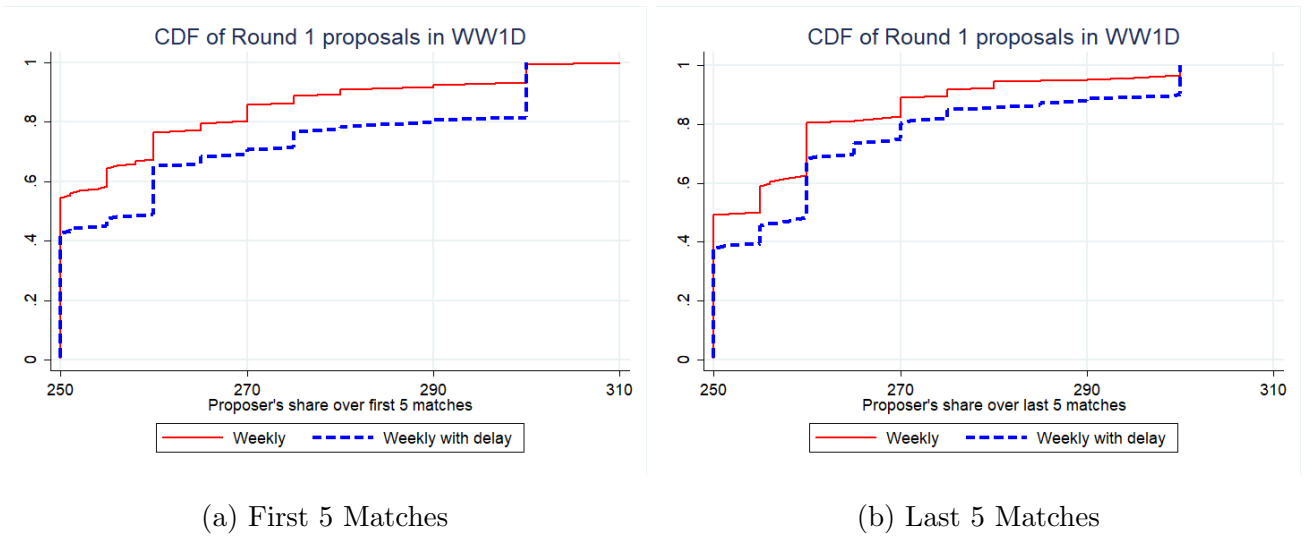


Figure 5: Round-1 Proposals in Treatment *WW1D*

Figure 5 shows the CDF of Round-1 proposals in Treatment *WW1D* by bargainer type. The solid line indicates the CDF for the weekly proposer (facing no front-end delay), and the dotted line indicates the CDF for the delayed weekly proposer. Whereas EXD predicts no difference, it is clear that the distribution of proposals by delayed weekly bargainers first-order stochastically dominates that of the weekly bargainers without front-end delay throughout, as alternatively predicted under QHD. This difference is highly significant both in the first 5 and the last 5 matches (KS test, p -value < 0.01 and p -value = 0.046, respectively). Hence, the front-end delay neutrality Prediction (4a) for Treatment *WW1D* is strongly rejected in favor of the alternative Prediction (4a') of a front-end delay advantage as under QHD (or any present bias). Instead of perceiving no difference in their bargaining power, the two types act in accordance with a shared understanding that a front-end delay confers an advantage due to present bias.

Result 5 (Front-End Delay Advantage in Treatment *WW1D*). *In Treatment WW1D, the initial proposals by delayed weekly bargainers significantly exceed and first-order stochastically dominate those by weekly bargainers without delay.*

5.2.2 Across Treatments

We now turn to testing the predictions of a basic delay advantage and front-end delay neutrality vs. advantage across treatments. Here, we generally exploit the feature of our design that we observe the same bargainer types against *different* opponent bargainer types. In contrast to the within-treatment comparisons, we therefore have to separately consider predictions for the “common type” as the initial proposer and as the initial respondent. Throughout, we focus on the data from the last 5 matches for these comparisons. Recall here that the rates of immediate agreement are similar across any two treatments.

First, we consider the basic delay advantage according to general Prediction (3b), which is common to both EXD and QHD: fixing the initial role, the weekly bargainer is predicted to obtain a greater share against a monthly bargainer (in Treatment *WM*) than against a delayed weekly bargainer (in Treatment *WW1D*). The weekly bargainer (with no delay) is the common type here, and we have to consider two comparisons, each comparing behavior across the two different opponent types (treatments): one for the weekly bargainer as the initial proposer and another for the weekly bargainer as the initial respondent. The two pairs of CDFs of initial proposals—by and to weekly bargainers, respectively—are shown in Figure 6.




Figure 6: Response to Different Types by Weekly – Last 5 Matches

Figure 6(a) compares Round-1 proposals by weekly bargainers to monthly bargainers, as in Treatment *WM* (solid), and to delayed weekly bargainers, as in Treatment *WW1D* (dashed), for the last 5 matches. Consistent with Prediction (3b), the former distribution clearly first-order stochastically dominates the latter. The difference is highly significant statistically (KS test, p -value < 0.01). This finding demonstrates that weekly bargainers perceive themselves in a much stronger initial proposing position against monthly bargainers than delayed weekly ones, in strong confirmation of Prediction (3b).³⁸

³⁸Notably, this finding contradicts near-future bias; see footnote 28.

Figure 6(b) compares Round-1 proposals to weekly bargainers by monthly bargainers, as in Treatment *WM* (solid), and by delayed weekly bargainers, as in Treatment *WW1D* (dashed), for the last 5 matches. Again, consistent with Prediction (3b), the former distribution quite obviously first-order stochastically dominates the latter, and this difference is also highly significant statistically (KS test, p -value < 0.01). What is remarkable here is that much of this difference is due to the fact that approximately 60% of monthly bargainers propose an equal split to a weekly bargainer, whereas only 40% of delayed weekly bargainers do so. This finding in turn demonstrates that, against the same weekly bargainer type, the monthly bargainers perceive their initial proposing position to be weaker than the delayed weekly ones perceive theirs, providing further strong confirmation of Prediction (3b).

Result 6 (Basic Delay Advantage across Treatments *WM* and *WW1D*). *The initial proposals by weekly bargainers to monthly bargainers in Treatment WM significantly exceed and first-order stochastically dominate those to delayed weekly bargainers in Treatment WW1D. The initial proposals to weekly bargainers by delayed weekly bargainers in Treatment WW1D significantly exceed and first-order stochastically dominate those by monthly bargainers in Treatment WM.*

Next, we turn to Prediction (3c) under EXD and its more permissive qualification (3c') under QHD, which concerns another instance of basic delay advantage. Here, the common type is the delayed weekly type, and we compare this type's outcomes against the delayed monthly type (Treatment *WM2D*) and against the weekly type with no delay (Treatment *WW1D*). Figure 7 provides the pairwise comparisons.




Figure 7: Response to Different Types by Delayed Weekly Bargainers— Last 5 Matches

Figure 7(a) compares Round-1 proposals by delayed weekly bargainers to delayed monthly bargainers, as in Treatment *WM2D* (solid), and to weekly bargainers (with no delay), as in Treatment *WW1D* (dashed). In this case, EXD predicts an unambiguously greater advantage against delayed monthly bargainers, whereas QHD implies that weekly bargainers with no delay may also be the weaker respondents if their present bias is sufficiently strong to outweigh their delay advantage in later rounds. The EXD Prediction (3c) is here rejected, since there is neither any qualitative first-order dominance relationship between the CDFs nor any statistically significant difference between them

(KS test, p -value = 0.349). In other words, delayed weekly bargainers perceive their initial proposing position to be roughly equally strong against both of these opponent types. This is consistent with Prediction (3c') under QHD (or a strong present bias).

Figure 7(b) compares Round-1 proposals made to delayed weekly bargainers by delayed monthly bargainers, as in Treatment *WM2D* (solid), and by weekly bargainers (with no delay), as in Treatment *WW1D* (dashed). Since an initial proposer's present bias is irrelevant to equilibrium, both EXD and QHD imply that weekly bargainers should claim more than delayed monthly ones from the same opponent type, here a delayed weekly bargainer. Our data reject this, however. There is a marginally significant difference with the opposite first-order stochastic dominance order (KS test, p -value = 0.09). Delayed monthly bargainers perceive their initial proposing position against delayed weekly ones to be somewhat stronger than weekly ones do, in contradiction to both Predictions (3c) and (3c').³⁹


Result 7 (Basic Delay Advantage across Treatments *WM2D* and *WW1D*). *The initial proposals by delayed weekly bargainers to delayed monthly bargainers in Treatment WM2D and those to weekly bargainers (with no delay) in Treatment WW1D are not significantly different in their distributions, and we observe no first-order stochastic dominance relationship. The initial proposals to delayed weekly bargainers by delayed monthly bargainers in Treatment WM2D marginally significantly exceed and first-order stochastically dominate those by weekly bargainers (with no delay) in Treatment WW1D.*

Finally, we turn to Predictions (4b) and (4b'), which concern front-end delay neutrality under EXD vs. front-end delay advantage under QHD (present bias). We investigate this by comparing Treatments *WM* and *WM2D*, which are strategically identical under EXD but not under QHD because a front-end delay lowers the initial respondent's cost of delay, thereby increasing bargaining power. Recalling that the initial proposer's present bias is irrelevant, the differential prediction under QHD is that Round-1 proposals by weekly bargainers (delayed or not) should claim a greater share in Treatment *WM* than in Treatment *WM2D*; a similar prediction is made for the initial proposals by monthly bargainers.

Figure 8(a) compares Round-1 proposals by weekly bargainers to monthly bargainers, as in Treatment *WM* (solid), and by delayed weekly bargainers to delayed monthly bargainers, as in Treatment *WM2D* (dashed). Under EXD, these are predicted to be the same, whereas under QHD, they should be larger in Treatment *WM* without front-end delay. The visual comparison suggests first-order stochastic dominance supporting the QHD prediction. However, there are somewhat more very high proposals (i.e., proposers claiming a very high share for themselves) in Treatment *WM2D*, and the difference in distributions is not statistically significant (KS test, p -value = 0.48). EXD's Prediction (4b) cannot explain the particular shape of CDFs yet cannot be rejected. Granting that our distributional test may be somewhat conservative, there is some, limited support for QHD's alternative Prediction (4b').

Figure 8(b) compares Round-1 proposals by monthly bargainers to weekly bargainers, as in Treatment *WM* (solid), and by delayed monthly bargainers to delayed weekly bargainers, as in Treatment

³⁹Near-future bias also leads to Prediction (3c), which is strongly rejected; see footnote 28.



(a) Weekly Proposers in *WM* v. *WM2D*

(b) Monthly Proposers in *WM* v. *WM2D*

Figure 8: Response to Different Types by Front-End Delay – Last 5 Matches

WM2D (dashed). In contrast to both Predictions (4b) and (4b'), the latter proposals first-order stochastically dominate the former, and this difference is highly statistically significant (KS test, p -value < 0.01). Delayed monthly bargainers claim *more* from delayed weekly ones than monthly bargainers do from weekly ones. Predictions (4b) and (4b') are overall rejected.

Result 8 (Front-End Delay Advantage Across Treatments *WM* and *WM2D*). *The initial proposals by weekly bargainers to monthly bargainers in Treatment WM tend to first-order stochastically dominate those by delayed weekly bargainers to delayed monthly bargainers in Treatment WM2D, but they are not significantly different statistically. The initial proposals by delayed monthly bargainers to delayed weekly bargainers in Treatment WM2D significantly exceed and first-order stochastically dominate those by weekly bargainers to monthly bargainers in Treatment WM.*

5.3 Summary and Discussion

Overall, we obtain rather strong support for the theoretical predictions under QHD (present bias). This is true above all in our within-treatment comparisons. The two failures—one part of the effective-delay advantage Prediction (3c') and one part of the front-end-delay advantage Prediction (4b')—occur in the naturally tougher comparisons *across* treatments. Remarkably, however, *both* occur with proposals by the same type in the same treatment: the delayed monthly bargainers in Treatment *WM2D* claim “too much” from their opponents, the delayed weekly bargainers; i.e., they perceive their position as “too strong.” Indeed, if we exclude all comparisons involving Treatment *WM2D*, thereby focusing on within-treatment Predictions (3a, *WM*) and (4a), as well as across-treatment Prediction (3b), then the behavioral support for the theory under QHD (present bias) is overwhelming. Our participants here convincingly demonstrate a shared understanding that a longer basic delay makes a bargainer weaker (both as the initial respondent and as the initial proposer), and a front-end delay makes a bargainer stronger (as the initial respondent).

What might explain then the theoretically surprising findings that, as the initial proposer, delayed monthly bargainers exert greater bargaining power against their delayed weekly opponents than (i) weekly bargainers (with no delay) against the same opponent types and (ii) monthly bargainers against weekly bargainers (neither with delay)? As pointed out in our discussion of the behavioral implications of other time preferences, hyperbolic discounting can permissively explain these features while predicting the same as QHD, and its predictions are strongly confirmed. From this perspective, we find evidence for present bias together with hyperbolic discounting. Still, QHD arguably captures the key features of behavior very well while being parametrically parsimonious and tractable.

As expected, we also find evidence for social preferences. Overall, more than 45% of proposals correspond to an equal split. There are, however, notable differences across treatments and bargainer types, as can be seen in the within-treatment comparisons. Indeed, the predictions’ confirmations in terms of first-order stochastic dominance all come with correspondingly sizable differences in the fraction of proposed equal splits. Relatedly, Treatment *WM2D* is the only one in which this fraction is almost identical for both types. The strong confirmation of Prediction (4a) via Treatment *WW1D*, a front-end delay advantage, could also be explained by social preferences for equalizing utility, given the treatment’s asymmetry in terms of the payoff timing of immediate agreements. However, the difference we find concerns the entire distribution and is substantial, which suggests that the contribution of social preferences is only partial. Confirming this interpretation, in both other treatments immediate agreements come with no payoff delay difference, and we still find a strong effect.

Overall, roughly 3 of 4 proposals are accepted immediately. While this means a great amount of efficiency, a fair number of proposals are still rejected. It seems clear that information is in fact incomplete because preferences are heterogeneous, and this incomplete information causes such rejections. However, our experimental procedure has arguably proven to be rather successful in establishing a common understanding of relative bargaining power based on time preferences, as intended. As part of this procedure, we made the different individual payoff delays within a match very salient. If some of the predictions’ confirmations appear to be “obvious” consequences of this, in spite of incomplete information about preferences (time, risk, social), then this is exactly the point to make: when two bargainers share a common understanding of who is more patient, then their behavior will reflect this in terms of greater bargaining power.⁴⁰

6 Concluding Remarks

Our findings confirm that patience is a source of bargaining power. We obtain this confirmation by contributing a novel between-subjects design that randomly manipulates time preferences at the individual level. As argued, we consider the behavioral success of the theory in our experiment—certainly relative to prior related studies—a consequence of how our design controls the inherently

⁴⁰This is also confirmed by our analysis of learning in Appendix C.

incomplete preference information that the participants and also we as researchers have. Since this issue arises importantly in many settings, we interpret this paper’s results also as a demonstration of our method’s wider usefulness.

This paper speaks to the large body of theoretical analyses of dynamic strategic interaction, where time preferences—with few exceptions, this means simply “the discount factor”—are a key driver of behavior. The general observation that time preferences matter in ways predicted by equilibrium, here established for the classic setting of indefinite alternating-offers bargaining, contributes a very positive message. It confirms a necessary condition for the empirical validity of basic theoretical exercises.

At the same time, beyond exponential discounting, patience becomes a more complicated notion; there is no longer simply “a” discount factor, but there are potentially many. We obtain strong evidence for a present bias, as parsimoniously captured by the quasi-hyperbolic discounting model. The fact that people seem to share a common understanding of this bias and strategically respond to it is encouraging and even calls for further theoretical analyses of dynamic strategic interaction with present-biased individuals.

At a more detailed level, however, we also obtain evidence for present bias as a feature of diminishing impatience, i.e., hyperbolic discounting. Our design offers a rare opportunity to investigate the strategic role of time preferences in such detail. While largely unexplored in strategic interaction (though see [Obara and Park, 2017](#), for a notable exception in the context of repeated games), our result may also inspire both empirical and theoretical future work in this direction. For instance, we suspect that diminishing impatience could also contribute towards explaining the frequently observed U-shaped agreement-time curves in bargaining settings with deadlines (i.e., disproportionately many agreements right at the beginning and just before the deadline) reported in the literature (e.g., [Roth, Murnighan, and Schoumaker, 1988](#); [Embrey, Fréchette, and Lehrer, 2014](#); [Karagözoğlu and Riedl, 2014](#); [Karagözoğlu, Keskin, and Özcan-Tok, 2019](#)).

References

- Akin, Zafer (2007), “Time inconsistency and learning in bargaining games.” *International Journal of Game Theory*, 36, 275–299.
- Andreoni, James and Charles Sprenger (2012), “Risk preferences are not time preferences.” *American Economic Review*, 102, 3357–3376.
- Augenblick, Ned, Muriel Niederle, and Charles Sprenger (2015), “Working over time: Dynamic inconsistency in real effort tasks.” *The Quarterly Journal of Economics*, 130, 1067–1115.
- Ausubel, Lawrence M., Peter Cramton, and Raymond J. Deneckere (2002), “Bargaining with incomplete information.” In *Handbook of Game Theory with Economic Applications*, volume 3, chapter 50, 1897–1945, Elsevier.
- Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak (1964), “Measuring utility by a single-response sequential method.” *Behavioral science*, 9, 226–232.
- Binmore, Ken (1987), “Perfect equilibria in bargaining models.” In *The economics of bargaining* (Ken Binmore and Partha Dasgupta, eds.), chapter 5, 77–105, Basil Blackwell.
- Binmore, Ken, Martin J. Osborne, and Ariel Rubinstein (1992), “Noncooperative models of bargaining.” In *Handbook of Game Theory with Economic Applications* (Robert Aumann and Sergiu Hart, eds.), volume 1, chapter 7, 179–225, North Holland.
- Binmore, Ken, Joseph Swierzbinski, and Chris Tomlinson (2007), “An experimental test of Rubinstein’s bargaining model.” Mimeo.
- Binmore, Kenneth G., Ariel Rubinstein, and Asher Wolinsky (1986), “The Nash bargaining solution in economic modelling.” *The RAND Journal of Economics*, 17, 176–188.
- Chakraborty, Anujit, Yoram Halevy, and Kota Saito (2020), “The relation between behavior under risk and over time.” *American Economic Review: Insights*, 2, 1–16.
- Coles, Melvyn G. and Abhinay Muthoo (2003), “Bargaining in a non-stationary environment.” *Journal of Economic Theory*, 109, 70–89.
- Dal Bó, Pedro and Guillaume R. Fréchette (2018), “On the determinants of cooperation in infinitely repeated games: A survey.” *Journal of Economic Literature*, 56, 60–114.
- Embrey, Matthew, Guillaume R. Fréchette, and Steven F. Lehrer (2014), “Bargaining and reputation: An experiment on bargaining in the presence of behavioural types.” *The Review of Economic Studies*, 82, 608–631.
- Fanning, Jack and Andrew Kloosterman (2019), “A simple experimental test of the coase conjecture: Fairness in dynamic bargaining.” Mimeo.

- Fischbacher, Urs (2007), “z-tree: Zurich toolbox for ready-made economic experiments.” *Experimental Economics*, 10, 171–178.
- Frederick, Shane, George Loewenstein, and Ted O’Donoghue (2002), “Time discounting and time preference: A critical review.” *Journal of Economic Literature*, 40, 351–401.
- Güth, Werner and Martin G. Kocher (2014), “More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature.” *Journal of Economic Behavior & Organization*, 108, 396–409.
- Güth, Werner, Rolf Schmittberger, and Bernd Schwarze (1982), “An experimental analysis of ultimatum bargaining.” *Journal of Economic Behavior and Organization*, 3, 367–388.
- Haan, Marco A. and Dominic Hauck (2019), “Games with possibly naive hyperbolic discounters.” Mimeo.
- Halevy, Yoram (2008), “Strotz meets Allais: Diminishing impatience and the certainty effect.” *American Economic Review*, 98, 1145–62.
- Halevy, Yoram (2015), “Time consistency: Stationarity and time invariance.” *Econometrica*, 83, 335–352.
- Ho, Kate and Robin S. Lee (2017), “Insurer competition in health care markets.” *Econometrica*, 85, 379–417.
- Karagözoğlu, Emin, Kerim Keskin, and Elif Özcan-Tok (2019), “Between anchors and aspirations: A new family of bargaining solutions.” *Review of Economic Design*, 23, 53–73.
- Karagözoğlu, Emin and Arno Riedl (2014), “Performance information, production uncertainty, and subjective entitlements in bargaining.” *Management Science*, 61, 2611–2626.
- Kim, Jeongbin (2019a), “Discounting, dynamic consistency, and cooperation in an infinitely repeated game experiment.” Mimeo.
- Kim, Jeongbin (2019b), “The effects of time preferences on cooperation: Experimental evidence from infinitely repeated games.” Mimeo.
- Laibson, David (1997), “Golden eggs and hyperbolic discounting.” *The Quarterly Journal of Economics*, 112, 443–478.
- Loewenstein, George and Drazen Prelec (1992), “Anomalies in intertemporal choice: Evidence and an interpretation.” *The Quarterly Journal of Economics*, 107, 573–597.
- Lu, Shih En (2016), “Self-control and bargaining.” *Journal of Economic Theory*, 165, 390–413.
- Manzini, Paola (2001), “Time preferences: Do they matter in bargaining?” Queen Mary University of London Working Paper No. 445.
- Muthoo, Abhinay (1999), *Bargaining Theory with Applications*. Cambridge University Press.
- Nash, John F. (1950), “The bargaining problem.” *Econometrica*, 155–162.

- Noor, Jawwad (2011), “Intertemporal choice and the magnitude effect.” *Games and Economic Behavior*, 72, 255–270.
- Obara, Ichiro and Jaeok Park (2017), “Repeated games with general discounting.” *Journal of Economic Theory*, 172, 348–375.
- O’Donoghue, Ted and Matthew Rabin (2015), “Present bias: Lessons learned and to be learned.” *American Economic Review: Papers & Proceedings*, 105, 273–279.
- Ok, Efe A. and Yusufcan Masatlioglu (2007), “A theory of (relative) discounting.” *Journal of Economic Theory*, 137, 214–245.
- Osborne, Martin J. and Ariel Rubinstein (1990a), *Bargaining and Markets*. Academic Press.
- Osborne, Martin J. and Ariel Rubinstein (1990b), *The Strategic Approach: A Model of Alternating Offers*, chapter 3, 29–68. Academic Press.
- Pan, Jinrui, Craig S. Webb, and Horst Zank (2015), “An extension of quasi-hyperbolic discounting to continuous time.” *Games and Economic Behavior*, 89, 43–55.
- Phelps, Edmund S. and Robert A. Pollak (1968), “On second-best national saving and game-equilibrium growth.” *The Review of Economic Studies*, 35, 185–199.
- Rapoport, Amnon, Eythan Weg, and Dan S. Felsenthal (1990), “Effects of fixed costs in two-person sequential bargaining.” *Theory and Decision*, 28, 47–71.
- Roth, Alvin E., J. Keith Murnighan, and Françoise Schoumaker (1988), “The deadline effect in bargaining: Some experimental evidence.” *The American Economic Review*, 78, 806–823.
- Rubinstein, Ariel (1982), “Perfect equilibrium in a bargaining model.” *Econometrica*, 50, 97–109.
- Schweighofer-Kodritsch, Sebastian (2018), “Time preferences and bargaining.” *Econometrica*, 86, 173–217.
- Sprenger, Charles (2015), “Judging experimental evidence on dynamic inconsistency.” *American Economic Review: Papers & Proceedings*, 105, 280–85.
- van Damme, Eric, Reinhard Selten, and Eyal Winter (1990), “Alternating bid bargaining with a smallest money unit.” *Games and Economic Behavior*, 2, 188–201.
- Weg, Eythan, Amnon Rapoport, and Dan S. Felsenthal (1990), “Two-person bargaining behavior in fixed discounting factors games with infinite horizon.” *Games and Economic Behavior*, 2, 76–95.
- Zwick, Rami, Amnon Rapoport, and John C. Howard (1992), “Two-person sequential bargaining behavior with exogenous breakdown.” *Theory and Decision*, 32, 241–268.

Appendices

A Proofs

Lemma 1

Proof. Define, for each player i , the function $f_i : [0, 1] \rightarrow [0, 1]$ as $f_i(U) = 1 - u_j^{-1}(U)$. If player j is the respondent and could obtain a fixed utility U by rejecting, then $1 - u_j^{-1}(U)$ is the maximal share of proposer i that j is willing to accept. Equation (3.1) then says that $x_n = f_{r_{n+1}}(\delta_{r_n}(n) \cdot u_{r_n}(x_{n+1}))$, whereby any sequence x_n corresponds to a history-independent equilibrium: in any round n , the proposing player offers share $1 - x_n$, thus keeping x_n for herself, and this is the smallest offer accepted by the responding player, who upon rejection would similarly capture x_{n+1} . (Note the indifference of the responding player, $u_{r_n}(1 - x_n) = \delta_{r_n}(n) \cdot u_{r_n}(x_{n+1})$.)

Take now any odd-numbered round N in which player 1 is the proposer, and consider the two extreme cases for responding player 2's continuation utility upon rejection: first, when it is minimal and equals zero, and second, when it is maximal and equals one. For each of these two cases, compute the implied backwards induction solution for the thus truncated game. Clearly, it has immediate agreement in every round, and starting from the respective extreme terminal values, it is characterized by the recursive equation (3.1) for all rounds up through round N . (The extreme shares $x_{N+1} = 0$ and $x_{N+1} = 1$ correspond to the extreme continuation utilities $U_2 = 0$ and $U_2 = 1$.) Define these two finite sequences as a_n^N and b_n^N , and—using assumption 3 with $\alpha \equiv \max\{\alpha_1, \alpha_2\}$ —observe that

$$\begin{aligned}
 |a_N^N - b_N^N| &= a_N^N - b_N^N \\
 &= f_1(0) - f_1(\delta_2(N)) \\
 &= u_2^{-1}(\delta_2(N)) - u_2^{-1}(0) \\
 &\leq \alpha \cdot \delta_2(N) \\
 |a_{N-1}^N - b_{N-1}^N| &= b_{N-1}^N - a_{N-1}^N \\
 &= f_2(\delta_1(N-1) \cdot u_1(f_1(\delta_2(N)))) - f_2(\delta_1(N-1) \cdot u_1(f_1(0))) \\
 &= u_1^{-1}(\delta_1(N-1) \cdot u_1(f_1(0))) - u_1^{-1}(\delta_1(N-1) \cdot u_1(f_1(\delta_2(N)))) \\
 &\leq \alpha \cdot (f_1(0) - f_1(\delta_2(N))) \\
 &\leq \alpha^2 \cdot \delta_2(N) \\
 &\vdots \\
 |a_1^N - b_1^N| &\leq \alpha^N \cdot \delta_2(N).
 \end{aligned}$$

Clearly, $|a_1^{2n-1} - b_1^{2n-1}| \rightarrow_{n \rightarrow \infty} 0$ (recall that we use only odd-numbered rounds), and hence $\lim_{n \rightarrow \infty} a_1^{2n-1} = \lim_{n \rightarrow \infty} b_1^{2n-1}$, which proves the claim, since $a_1^{2n-1} \geq x_1 \geq b_1^{2n-1}$ for all n . \square

Proposition 1

Proof. Consider any odd-numbered round N in which player 1 is the proposer, and suppose the supremal equilibrium continuation utility of player 2 takes the highest possible value of 1. Then there exists an equilibrium with the outcome that players agree in round 1 and proposing player 1 obtains share a_1^N , defined in the proof of Lemma 1. Similarly, supposing the infimal equilibrium continuation utility of player 2 takes the lowest possible value of 0, there exists an equilibrium with the outcome that players agree in round 1 and proposing player 1 obtains share b_1^N , defined in the proof of Lemma 1. Now, any equilibrium utility value U_1 of player 1 (as of round 1) satisfies $u_1(a_1)^N \geq U_1 \geq u_1(b_1^N)$, whereby Lemma 1 proves its uniqueness. A similar argument proves the uniqueness of player 2's equilibrium utility. Both are uniquely obtained in the immediate-agreement equilibrium characterized by the sequence of Lemma 1. \square

Prediction 1

Proof. Immediate agreement (1) is a general implication, irrespective of how players discount utility. The other predictions require proof. In preparation, note that defining $f(U) \equiv 1 - u^{-1}(U)$ for any $U \in [0, 1]$, Proposition 1 implies that the unique equilibrium is characterized by

$$x_1^E = f(\phi_2 \delta u(f(\phi_1 \delta u(x_1^E)))) \text{ and } x_2^E = f(\phi_1 \delta u(x_1^E)), \quad (\text{A.1})$$

where x_i^E is the share that individual i obtains in immediate agreement whenever she gets to propose. This share x_i^E obtains as the unique (and interior) fixed point of the function $g_i(q) \equiv f(\phi_j \delta u(f(\phi_i \delta u(q))))$, defined for any $q \in [0, 1]$.⁴¹ The characterization covers all matches of all treatments.

For the Proposer Advantage (2), simply observe that $x_i^E > u^{-1}(\phi_i \delta u(x_i^E)) = 1 - f(\phi_i \delta u(x_i^E)) = 1 - x_j^E$.

For the Basic Delay Advantage (3), observe that $\phi_1 > \phi_2$ implies $g_1(q) > g_2(q)$ for all $q \in [0, 1]$, and therefore $x_1^E > x_2^E$ (comparison of proposer shares), which is equivalent to $1 - x_2^E > 1 - x_1^E$ (comparison of respondent shares). Given A.1, this covers all parts.

Finally, (4) follows directly from the irrelevance of front-end delay under exponential discounting, as explained earlier. \square

Prediction 2

Proof. Again, (1) requires no specific proof. For the remainder, note that the second-round continuation equilibrium is characterized by the shares x_i^E solving the two equations (A.1). Backwards induction then yields immediate agreement in the first round, with the initial proposer's share given by

$$x_1^Q = f(\beta_2 \phi_2 \delta u(x_2^E)).$$

⁴¹Our preference assumptions imply that each g_i is continuous and increasing from $g_i(0) > 0$ through $g_i(1) < 1$, whereby a fixed point exists and any fixed point is interior. Moreover, by our third preference assumption, each g_i has a slope less than one, so there is a unique fixed point.

(2) follows straight from the corresponding proof for EXD upon noting that $\beta_2 \leq 1$ implies $x_1^Q \geq x_1^E$, since $x_1^E > 1 - x_2^E = u^{-1}(\phi_1 \delta u(x_1^E)) \geq u^{-1}(\beta_1 \phi_1 \delta u(x_1^E))$.

For (3a), first observe that *WM* has $\beta_1 = \beta_2 = \beta$ and that the respondent's continuation share is smaller for the monthly than the weekly bargainer from EXD. Hence, the initial proposer's share x_1^Q is greater (equivalently, the initial respondent's share $1 - x_1^Q$ is smaller) when the weekly bargainer initially proposes against the monthly bargainer than when the monthly bargainer initially proposes against the weekly bargainer. Second, observe that *WM2D* has $\beta_1 = \beta_2 = 1$, whereby predictions are as under EXD.

For (3b), observe that the weekly bargainer's continuation share is greater against the monthly bargainer (*WM*) than against the delayed weekly bargainer (*WW1D*), both as the initial proposer and as the initial respondent, from EXD. Hence, when the weekly bargainer is the initial respondent, $(\phi_2, \beta_2) = (1, \beta)$, $1 - x_1^Q$ is greater against the monthly bargainer, $(\phi_1, \beta_1) = (\phi, \beta)$, than against the delayed weekly bargainer, $(\phi_1, \beta_1) = (1, 1)$. When the weekly bargainer is the initial proposer, a responding delayed weekly bargainer is unaffected by present bias, whereas a responding monthly bargainer is additionally weakened by it; this implication follows also for the cross-treatment comparison of the weekly bargainer's shares as the initial proposer.

For (3c'), first observe that with the initial respondent's type equal to $(\phi_2, \beta_2) = (1, 1)$, her continuation share—hence also $1 - x_1^Q$ —is smaller against the weekly than the monthly bargainer, as under EXD. Second, fixing $(\phi_1, \beta_1) = (1, 1)$, it should be clear from continuity that a violation of the prediction under EXD—meaning x_1^Q is smaller when $(\phi_2, \beta_2) = (\phi, 1)$ than when $(\phi_2, \beta_2) = (1, \beta)$ —is obtained as ϕ approaches one while β approaches zero.

For (4a'), observe that when the weekly bargainer is the initial proposer, then $x_1^Q = x_1^E$, while when the weekly bargainer is the initial respondent, then $x_1^Q > x_1^E$.

For (4b'), observe that in *WM*, under either assignment of roles, $x_1^Q > x_1^E$, whereas in *WM2D*, under either assignment of roles, $x_1^Q = x_1^E$. □

B Additional Figures

Efficiency




Figure 9: The Proportions of Agreements over Rounds – First 5 Matches




Figure 10: The Proportions of Immediate Agreements – All Matches

Proposer Advantage: First 5 Matches' Proposals and Accepted Proposals




Figure 11: Proposer Advantage – First 5 Matches' Proposals

Proposer Advantage: Final Payoffs incl. Random Termination




Figure 12: Final Payoffs (All) – First and Last 5 Matches

Proposer Advantage: Final Payoffs excl. Random Termination




Figure 13: Final Payoffs (excl. Random Terminations) – First and Last 5 Matches

Basic Delay and Front-End Delay Advantages: Proposals over Matches




Figure 14: Round-1 Proposals over Matches in Treatment *WM*




Figure 15: Round-1 Proposals over Matches in Treatment *WM2D*






Figure 16: Round-1 Proposals over Matches in Treatment *WW1D*

Basic Delay and Front-End Delay Advantages: Accepted Proposals




(a) First 5 Matches




(b) Last 5 Matches

Figure 17: Accepted Proposals in Treatment *WM* – First and Last 5 Matches




(a) First 5 Matches




(b) Last 5 Matches

Figure 18: Accepted Proposals in Treatment *WM2D* – First and Last 5 Matches



(a) First 5 Matches



(b) Last 5 Matches

Figure 19: Accepted Proposals in Treatment *WW1D* – First and Last 5 Matches

C Learning

We explore here whether and how a proposer’s behavior is affected by the bargaining outcome in the previous match. Table 2 reports the regression results in which the dependent variable is the proposer’s proposed (own) share in Round 1 of each match (excluding the very first one, which has no history).

Table 2: Learning

	Accepted proposals		All proposals in Round 1	
	(1) Match 2-5	(2) Match 6-10	(3) Match 2-5	(4) Match 6-10
Match	-2.753** (1.111)	-1.500** (0.690)	-2.960** (1.055)	-1.443* (0.695)
Accepted Share	0.221** (0.101)	0.393** (0.140)		
Accepted Round	3.904 (2.460)	5.959** (2.198)		
Round 1 Share			0.010 (0.036)	0.042 (0.056)
Round 1 Accept			-41.723** (16.823)	-104.134** (36.187)
Round 1 Accept \times Round 1 Share			0.160* (0.080)	0.338** (0.126)
Constant	215.125*** (27.188)	169.088*** (32.635)	275.465*** (5.721)	283.736*** (7.227)
R-squared	0.043	0.102	0.027	0.074
Observations	615	796	696	870

Notes: OLS regression with the proposer’s proposed (own) share in Round 1 as the dependent variable. The independent variable “Accepted Share” is one’s own share that was accepted in Round 1 of the previous match. “Accepted Round” represents the number of the round in which the proposal was accepted in the previous match. “Round 1 Share” is one’s own share that was proposed in Round 1 of the previous match. “Round 1 Accept” is a dummy variable that takes a value of 1 if and only if the proposal was accepted in Round 1 of the previous match. Clustered standard errors at the session level are reported in parentheses. ***Significant at 1%; **5%; *10%

Overall, our participants were more likely to offer higher shares to respondents as they gained experience. In columns (1) and (2), we restrict our attention to proposals that were accepted in the previous match. A higher share in the previous match meant that a participant became more likely to claim a higher share as the proposer in Round 1 of the subsequent match. The magnitude of this association became stronger in the later part of the experiment. An agreement reached only in later rounds also made the proposals in Round 1 more aggressive, but this was significant only in the later matches.

We also look at all proposals in the previous match in columns (3) and (4), with a particular focus on whether proposals were accepted in Round 1. The fact that a proposal was rejected in Round 1 of the previous match led proposers to make less aggressive proposals in the subsequent match. Interestingly, the tendency of higher shares in the previous match to lead to more aggressive proposals was significantly stronger for accepted proposals than for rejected proposals. In summary, participants’ learning from their past experiences made them behave more consistently with the theoretical predictions.

D Experimental Instructions - Treatment *WM*

Welcome to the experiment. Please read these instructions carefully; the payment you will receive from this experiment depends on the decisions you make. The amount you earn will be paid through **VENMO**.

Your Payment Type and Match

At the beginning of the experiment, one-half of the participants will be randomly assigned to be Payment **Type A** and the other half to be Payment **Type B**. Your payment type will remain fixed throughout the experiment. Your payment type will affect when you will be paid, which will be explained below.

The experiment consists of 10 **matches**. At the beginning of each match, one Type A participant and one Type B participant are randomly paired. The pair is fixed **within the match**. After each match, participants will be randomly repaired, and new pairs will be formed. You will not learn the identity of the participant you are paired with, nor will that participant learn your identity—even after the end of the experiment.

Your Decisions in Each Match

Round 1: At the beginning of Round 1, one participant will be randomly assigned to the role of a **proposer** and the other participant to the role of a **responder**. Each participant in a match has 50-50 chance to be the proposer and to be the responder regardless of his/her payment type.

The proposer is then asked to propose how to split 500 tokens (= \$50) between the two participants as:

“_____ tokens for yourself and _____ tokens for the other person.”

After observing the split proposed by the proposer, the responder decides whether to accept or reject the proposed split.

Outcome, Termination, and Transition to Next Round: The outcome of Round 1 depends on whether the split proposed by the proposer is accepted or rejected.

1. If the responder **accepts** the proposed split, both participants will receive the amounts of tokens as proposed, and the match will be terminated.
2. If the responder **rejects** the proposed split, then the match will proceed to the next round with 75% (3/4) chance or be terminated with 25% (1/4) chance. This is as if we were to roll a 100-sided die and continue if the selected number is less than or equal to 75 and end if the number chosen is larger than 75.
 - (a) If a match is **terminated** after a rejection of a proposed split, both participants will receive 0 tokens for the match.
 - (b) If the match **proceeds** to the next round, then the proposer-responder roles are alternated. That is, the participant who is the proposer in the current round will become the responder in the next round, and vice versa. The number of tokens the participants receive will be determined by the outcome of the subsequent rounds.

Round $K > 1$: In Round $K > 1$, the participant who was the proposer in Round $(K - 1)$ becomes the responder, and the participant who was the responder in Round $(K - 1)$ becomes the proposer. The proposer is then asked to propose how to split 500 tokens (= \$50) between the two participants. After observing the split proposed by the proposer, the responder decides whether to accept or reject the proposed split.

The rest of the procedures determining the outcome, termination of the round, and transition to next round, is the same as those in Round 1.

Information Feedback

- At the end of each **round**, you will be informed about the proposal made by the proposer and the accept/reject decision made by the responder.
- At the end of each **match**, you will be informed **when and how much** you are going to be paid.

Your Monetary Payments

At the end of the experiment, one match out of 10 will be randomly selected for your payment. Every match has an equal chance to be selected for your payment so that it is in your best interest to take each match seriously. Participants will receive the amounts of tokens according to the outcome from the selected match with the exchange rate of 1 token = \$0.1.

When you are going to be paid depends on (1) your payment type and (2) the round in which the proposed split is accepted.

If you are **Type A**, you may be paid today or in a few **weeks**. If a proposed split is accepted in Round 1, you will be paid today right after the experiment. If a proposed split is accepted in Round 2, you will be paid in one week. If a proposed split is accepted in Round $K > 1$, you will be paid in $(K - 1)$ weeks.

If you are **Type B**, you may be paid today or in a few **months**. If a proposed split is accepted in Round 1, you will be today right after the experiment. If a proposed split is accepted in Round 2, you will be paid in one month. If a proposed split is accepted in Round $K > 1$, you will be paid in $(K - 1)$ months.

The following table summarizes the schedule of payment for each type:

If a proposed split is accepted in	Type A will be paid	Type B will be paid
Round 1	Today	Today
Round 2	In 1 week	In 1 month
Round 3	In 2 weeks	In 2 months
Round 4	In 3 weeks	In 3 months
Round 5	In 4 weeks	In 4 months
.....
Round K	In $(K - 1)$ weeks	In $(K - 1)$ months

Any amount you are supposed to receive will be paid electronically via VENMO.

In addition to your earnings from the selected match, you will receive a **show-up fee of \$10** through VENMO, right after the experiment.

A Practice Match

To ensure your comprehension of the instructions, you will participate in a practice match. The practice match is part of the instructions and is not relevant to your cash payment; its objective is to get you familiar with the computer interface and the flow of the decisions in each round of a match. Once the practice match is over, the computer will tell you “The official matches begin now!”

Rundown of the Study

1. At the beginning of the experiment, your payment type will be randomly determined. Your payment type will remain fixed throughout the experiment.
2. At the beginning of each match, one Type A participant and one Type B participant are randomly paired.
3. At the beginning of Round 1, one participant will be randomly assigned to the role of a proposer and the other to the role of a responder.
4. The proposer then proposes how to split 500 tokens (= \$50).
5. If the responder accepts the proposed split, both participants will receive the amounts of tokens as proposed, and the match will be terminated.
6. If the responder rejects the proposed split, then the match will proceed to the next round with 75% ($3/4$) chance or be terminated with 25% ($1/4$) chance. If a match is terminated after the rejection of a proposed split, both participants will receive 0 tokens for the match.
7. If the match proceeds to the next round, then the proposer-responder roles are alternated.
8. At the end of the experiment, one of 10 matches will be randomly selected for payment. For the selected match, the timing of your payment depends on (1) your payment type and (2) the round in which the proposed split was accepted. All your earnings will be paid to you through VENMO.
9. For Type A, you may be paid today or in a few weeks. For Type B, you may be paid today or in a few months.
10. In addition to your earnings from the selected match, you will receive a show-up fee of \$10 right after the experiment.

Administration

Your decisions, as well as your monetary payment, will be kept confidential. Remember that you have to make your decisions entirely on your own; please do not discuss your decisions with any other participants. Upon finishing the experiment, you will be asked to sign your name to acknowledge your receipt of the payment. You are then free to leave. If you have any question, please raise your hand now. We will answer your question individually.

E Selected z-Tree Screen-shots

Timing of Payment		
Round	The Other's Type (A)	Your Type (B)
Round 1	Today	Today
Round 2	In 1 week	In 1 month
Round 3	In 2 weeks	In 2 months
...

Previous Round Summary will be provided here.

Match: 1 **Round: 1**

Your Role in this round = **PROPOSER**.

Payment timing for the accepted split:
 Your tokens **Today**.
 The other person's tokens **Today**.

Please propose a split of 500 tokens:

Tokens for yourself:

Tokens for the other person:

Figure 20: *Proposer's* Screen

Timing of Payment		
Round	Your Type (A)	The Other's Type (B)
Round 1	Today	Today
Round 2	In 1 week	In 1 month
Round 3	In 2 weeks	In 2 months
...

Previous Round Summary will be provided here.

Match: 1 **Round: 1**

Your Role in this round = **RESPONDER**.

Payment timing for the accepted split:
 Your tokens **Today**.
 The other person's tokens **Today**.

The proposed split from the other person:
 250 for you and 250 for the other.

Would you like to accept / reject the proposal?
Please select a column

Accept

Reject

Figure 21: *Responder's* Screen

F Elicited Time Preferences and Behavior

We also elicited conventional measures of time preferences from our participants. This served two purposes: First, we can thereby test whether the random assignment to treatment and also bargainer type was indeed successful with regards to the underlying time preferences, and second, we can also relate those conventional measures to behavior, as a complement to our main analysis.

Elicitation Procedure. We administered our elicitation task in only 4 out of the 6 sessions in each treatment (228 out of 348 participants), where it followed the bargaining games. Participants were not informed about this elicitation task beforehand, and they received all payoff-relevant information from their choices only at the very end of the experiment. The elicitation task asked participants to make 8 blocks of binary decisions between a sooner payment (option A) and a later payment (option B). In each block, one of the two was a fixed amount (either \$4 or \$10), and the other amount increased from \$0.01 in minimal steps of \$0.01 to \$10.00, resulting in effectively 1,000 binary decisions (rows) per block. Participants were asked for their switching point in terms of the varying option’s amount, which they had to enter. The computer would automatically select the fixed option in all rows with a smaller varying amount and the varying option in all rows with a larger such amount. One row would be selected at random and the decision implemented, for one randomly drawn block. In essence, this is a version of the BDM ([Becker, DeGroot, and Marschak, 1964](#)) method, hence incentive compatible, but explained via a price list. The full instructions and a screenshot are available at the end of this section.

Table 3: Description of the Elicitation Task

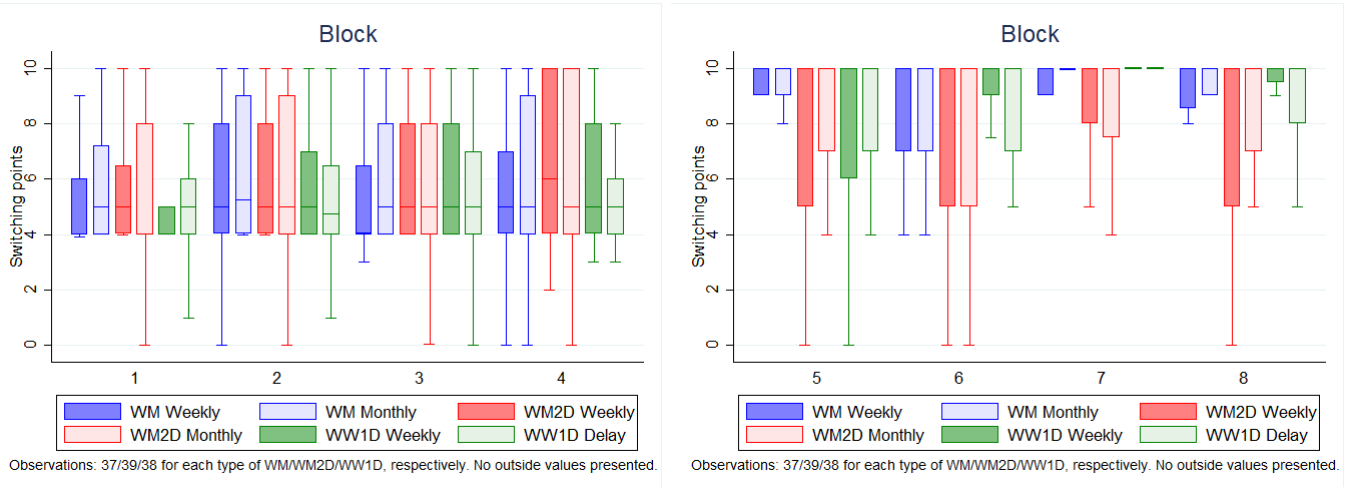
Switching	Sooner \Rightarrow Later			
Block	(1)	(2)	(3)	(4)
Sooner	\$4 Today	\$4 Today	\$4 1 month	\$4 1 month
Later	\$X 1 week	\$X 1 month	\$X 1 month and 1 week	\$X 2 months
Switching	Sooner \Leftarrow Later			
Block	(5)	(6)	(7)	(8)
Sooner	\$X Today	\$X Today	\$X 1 month	\$X 1 month
Later	\$10 1 week	\$10 1 month	\$10 1 month and 1 week	\$10 2 months

*Note: X denotes the amounts that vary from 0.01 to 10.

Table 3 provides an overview of the details of the task. The block numbers correspond to their order in

the task. There were four different sooner and later payment combinations: (1) sooner payment today and later payment in 1 week, (2) sooner payment today and later payment in 1 month, (3) sooner payment in 1 month and later payment in 1 month plus 1 week, and (4) sooner payment in 1 month and later payment in 2 months. For the first 4 blocks, the sooner payment was fixed at \$4.00 while the later payment ranged from \$0.01 to \$10.00. For the last 4 blocks, the later payment was fixed at \$10.00, and the sooner payment ranged from \$0.01 to \$10.00.

Distributions of Switching Points. We first compare the distributions of switching points X_k , where $k \in \{1, 2, \dots, 8\}$ refers to the block number, by treatment and bargainer type, to check whether our randomization in terms of underlying time preferences was successful. Figure 22 provides the corresponding box plots. We use the same test as for our bargaining predictions, the Kolmogorov-Smirnov test, to compare the switching point distributions on all 8 blocks. Since we test bargaining predictions both concerning comparisons between the two bargainer types within any treatment and between treatments for a given bargainer type, we carry out analogous tests on the time preference task responses. Comparing, first, the switching points between the two bargainer types within any treatment—e.g., weekly vs. monthly in treatment *WM*—we find no significant differences (8 binary comparisons per treatment times 3 treatments, hence 24 binary comparisons, all p -values greater than 0.239). Second, and given this finding, we compare responses between various pairs of treatments—e.g., *WM* vs. *WM2D*—with a similar result (8 binary comparisons per treatment pairing times 3 treatment pairings, hence 24 binary comparisons, all p -values greater than 0.226).⁴² Overall, we therefore conclude that our randomization into treatments and types in terms of underlying time preferences was successful indeed.



(a) Blocks 1-4

(b) Blocks 5-8

Figure 22: Distribution of Switching Points by Type/Treatment/Block

⁴²We run the same test for weekly types only, where there are three treatment comparisons (there are weekly types in all treatments) and for monthly types only, where there is one treatment comparison (*WM* vs. *WM2D*). This results in $(3 + 1) \cdot 8 = 32$ binary comparisons, and all except three of them have p -values greater than 0.375. The smallest three equal 0.117, 0.123 and 0.167, so may be considered borderline. However, all of them concern comparisons of weekly types for trade-offs with a month's delay, namely X_4 and X_8 , which are not the relevant ones for their bargaining.

Relation to Bargaining Behavior. We next relate our elicitation to bargaining behavior. The elicitation task is designed to infer parameters of (β, δ) -discounting, under the assumption that the participants are approximately risk neutral together with the standard narrow bracketing assumption (recall here the small stakes of at most \$10). We first estimate these for every participant, using the switching points for indifference equations—e.g., $4 = \beta\delta X_1$ and $4 = \delta X_3$, or $X_5 = \beta\delta 10$ and $X_7 = \delta 10$; details below—and then relate proposer as well as respondent behavior to the parameter estimates using regressions.

To estimate the two parameters we use for each participant the responses to the blocks that correspond to their type in bargaining; i.e., for a weekly type we consider X_1, X_3, X_5 , and X_7 , which involve a delay of a week, and for a monthly type we consider the other four, which involve a delay of a month (in both cases irrespective of whether their type has front-end delay). We then exclude participants whose responses on their relevant subset are inconsistent or do not allow us to infer indifference.⁴³ For the remaining participants, we compute (β, δ) once from the two relevant sooner-to-later switching points among the first four blocks and again from the two relevant later-to-sooner switching point among the last four blocks, and we then take the average of the two for each parameter to reduce measurement error. For instance, for a weekly type, we compute δ_w as the average of $\delta_{w(1)} = 4/X_3$ and $\delta_{w(2)} = X_7/10$, and then β_w as the average of $\beta_{w(1)} = 4/\delta_{w(1)}X_1 = X_3/X_1$ and $\beta_{w(2)} = X_5/\delta_{w(2)}10 = X_5/X_7$; similarly, for a monthly type, where we denote estimates by (β_m, δ_m) . The results are summarized in Table 4 in terms of averages with standard deviations, and in Figure 23 in terms of box-plots, by types and treatments.

Table 4: Average Elicited Time Preferences by Type

	Treatments					
	WM		WM2D		WW1D	
	Weekly	Monthly	Weekly	Monthly	Delay	Weekly
β_w	1.08 (0.42)	n.a.	0.97 (0.11)	n.a.	0.99 (0.08)	1.02 (0.11)
δ_w	0.85 (0.19)	n.a.	0.88 (0.13)	n.a.	0.88 (0.13)	0.86 (0.14)
β_m	n.a.	1.00 (0.13)	n.a.	0.99 (0.15)	n.a.	n.a.
δ_m	n.a.	0.85 (0.17)	n.a.	0.90 (0.13)	n.a.	n.a.
Obs.	26	25	26	18	27	28
# excluded	11	12	13	21	11	10

*Note: Standard deviations in parentheses.

Table 4 shows that around a third of participants per type and treatment had to be excluded, with the exception of monthly bargainers in Treatment *WM2D*, where this was one half. The average β is very similar in

⁴³Inconsistency refers to assumed impatience and transitivity. It means here (i) $X_k < 4$, or (ii) $X_k = 10$ and $X_{k+4} > 4$, for at least one of the two relevant $k \in \{1, 2, 3, 4\}$; moreover, while $X_k = 10$ together with $X_{k+4} < 4$ is not inconsistent, it does not allow to establish indifference because a highly impatient person may *strictly* prefer the fixed sooner amount of \$4 in block k over the maximal possible switching point of \$10.

all six cases, ranging from 0.97 to 1.08, which supports the quasi-hyperbolic structure. Moreover, the standard deviations are of similar sizes, except for weekly types in Treatment *WM*, where there are two outliers. Also the average δ is very similar in all six cases, ranging from 0.85 to 0.90, and all standard deviations are of similar size. There is no tendency for δ_m to be smaller than δ_w , however, even though it is based on a month's delay as opposed to a week's delay.

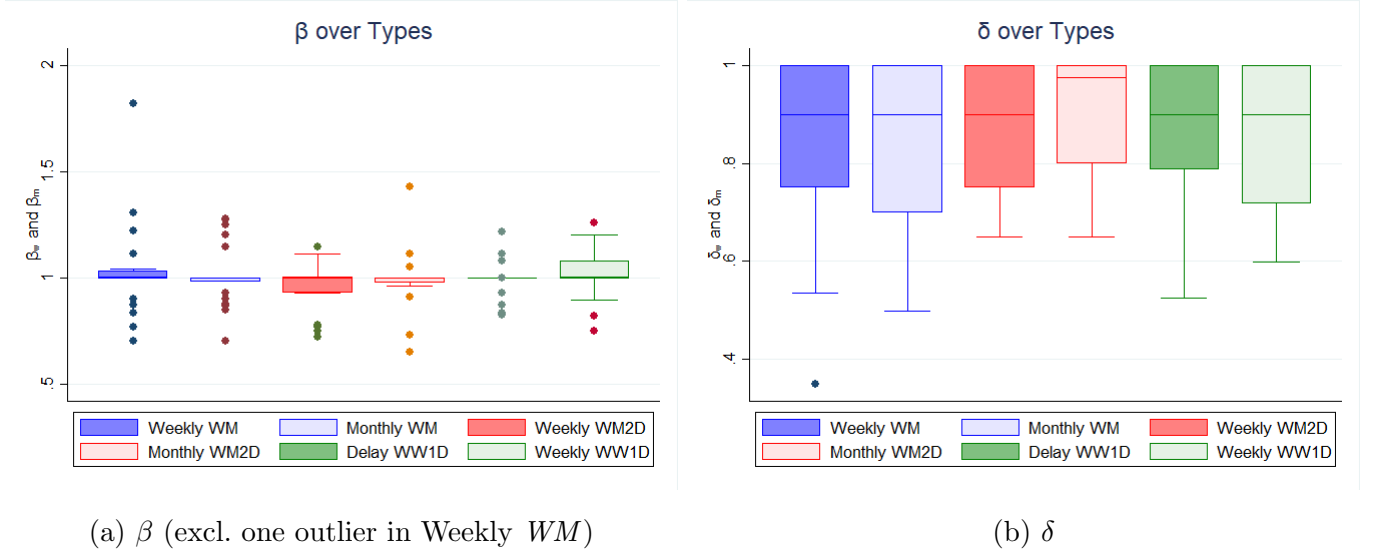


Figure 23: Distribution of β and δ by Type/Treatment

Figure 23 presents the underlying distributional information as box-plots, also including outside values. (One weekly type in Treatment *WM* has a far outlying β_w close to 3, which is not shown.) The median values of β are all equal to one, and most of the mass lies around one in all cases, so the distributions are rather similar. The median values of δ are equal to 0.90 in all cases except for the monthly types of Treatment *WM2D* for whom the median equals 0.97, and the distributions are quite similar too (again, regardless of whether the delay is a week or a month).

We now use our estimates as regressors in two specifications, one regarding proposer behavior and another regarding respondent behavior. Table 5 presents the results of OLS regressions of Round-1 proposals of various types in the treatments on the proposer's time preference parameters (and a constant). The explanatory power of these parameters is extremely weak, only in two cases is there a weakly statistically significant (partial) correlation with the proposals made (once positive for β in specification 3, and once positive for δ in specification 5), and there is no discernible pattern in terms of coefficient signs. When pooling all observations (specification 7), we observe what may be the expected positive signs for both parameters, though the coefficient on β is estimated very close to zero and far from being significantly different, and the co-efficient on δ is barely significant (at the 10%-level); the explanatory power is once again extremely weak.⁴⁴

⁴⁴If we only include observations where the proposer demands strictly more than half the cake (> 250) in the pooled specification (352 observations), then both positive co-efficients become greater in size, that on β turns significant at the 10%-level, and that on δ turns significant at the 1%-level. However, they still explain very little of the variation in proposals. Moreover, in the separate regressions per type/treatment, the results remain very similar also under that sample restriction.

Table 5: β , δ , and Round-1 Proposer behavior (OLS)

	WM		WM2D		WW1D		ALL
	(1) Weekly	(2) Monthly	(3) Weekly	(4) Monthly	(5) Delay	(6) Weekly	(7)
β	26.15 (18.44)	-17.36 (29.11)	61.57* (23.66)	-14.03 (13.85)	-11.89 (18.88)	-29.16 (13.65)	5.95 (5.60)
δ	82.61 (52.43)	-18.08 (23.29)	-28.06 (17.77)	0.98 (25.00)	33.87* (12.50)	28.31 (28.51)	18.67* (9.08)
Constant	170.7* (66.37)	286.3*** (38.90)	231.2*** (10.99)	281.2*** (22.19)	245.1*** (24.86)	265.2*** (24.09)	240.9*** (11.78)
R^2	0.070	0.008	0.076	0.006	0.029	0.032	0.006
Obs.	132	121	132	76	123	155	739

Notes: Dependent variable: Proposer's Round-1 share. Clustered standard errors at the session level in parentheses.

*** Significant at the 1%-level.

** Significant at the 5%-level.

* Significant at the 10%-level.

Table 6 presents the results of Probit regressions of Round-1 respondent behavior (acceptance=1, rejection=0) of various types in the treatments on the share offered together with the two time preference parameters (and a constant). The results mirror those for proposer behavior: The time preference measures are also hardly related to respondent behavior.

Table 6: β , δ , and Round-1 Respondent behavior (Probit)

	WM		WM2D		WW1D		ALL
	(1) Weekly	(2) Monthly	(3) Weekly	(4) Monthly	(5) Delay	(6) Weekly	(7)
Own share	0.054*** (0.007)	0.048*** (0.006)	0.027** (0.012)	0.042*** (0.011)	0.067*** (0.006)	0.050*** (0.007)	0.041*** (0.005)
β	-0.656 (0.607)	1.725 (3.711)	-0.524 (1.765)	-0.095 (0.565)	-2.335** (1.017)	-0.286 (2.464)	-0.054 (0.221)
δ	-2.531 (1.841)	-3.445*** (1.266)	0.696 (1.534)	-0.801 (1.927)	-0.281 (1.648)	2.419* (1.434)	-0.643 (0.685)
Constant	-9.235*** (2.587)	-9.000** (4.105)	-5.915 (4.289)	-8.180* (4.431)	-12.72*** (2.722)	-13.14*** (3.403)	-8.395*** (1.443)
Pseudo- R^2	0.316	0.477	0.174	0.322	0.374	0.326	0.229
Obs.	128	129	128	104	147	125	761

Notes: Dependent variable: Respondent's Round-1 acceptance=1, rejection=0. Clustered standard errors at the session level in parentheses.

*** Significant at the 1%-level.

** Significant at the 5%-level.

* Significant at the 10%-level.

Overall, even within a given treatment for a given bargainer type, conventional time preference measures

hardly explain behavior. Potential reasons include behaviorally relevant confounds (social preferences and risk attitudes, belief formation about the opponent) or also a relatively low signal-to-noise ratio of such measures. As such, the findings lend further support to our study's design and analysis.

Instructions for Elicitation Task and Selected z-Tree Screen-shot.

Instructions

In this task, we will ask you to make decisions for 8 blocks of questions. In each block, there are 1,000 questions. For each question, you can choose one of two options - Option A, which pays you sooner, and Option B, which pays you later.

After you answer all questions, one question will be randomly selected and the option you chose on that question will determine your earnings. Each question is equally likely to be chosen for payment. Obviously, you have no reason to misreport your preferred choice for any question, because if that question gets chosen for payment, then you would end up with the option you like less.

For example, the questions in one block are as follows. Note that each row corresponds to a question so that you have to choose one option for each row.

Questions	Option A Today	Option B in 1 month
1	\$4.00	\$0.01
2	\$4.00	\$0.02
3	\$4.00	\$0.03
⋮	⋮	⋮
999	\$4.00	\$9.99
1,000	\$4.00	\$10.00

It is natural to expect that you will choose Option A for at least the first few questions, but at some point switch to choosing Option B. In order to save time, you can report at which dollar value of Option B you'd switch. The computer program can then 'fill out' your answers to all 1,000 questions based on your reported switching point (choosing Option A for all questions before your switching point, and Option B for all questions at and after your switching point).

Timing of payment: The 8 blocks will differ in the following two ways: (1) the timings of sooner and later payments:

- Between payment **today** and payment in **1 week**.
- Between payment **today** and payment in **1 month**.
- Between payment in **1 month** and payment in **1 month and 1 week**.
- Between payment in **1 month** and payment in **2 months**.

and (2) whether you are asked to switch from Option A to Option B, or from Option B to Option A.

Payment: At the end of the experiment, one question in one of the blocks will be randomly selected for payment. The selected question and the block as well as your choice for the question will be displayed on your screen. Then the payment will be made on the designated date through VENMO. For example, 1. If your choice in the randomly selected question was to receive a payment today, then you will be paid through VENMO right after the experiment. 2. If your choice in the randomly selected question was to receive a payment in the future, you will be paid on the designated date through VENMO.

Rundown of the Study

- 1. There are 8 blocks of questions, each of which you will be asked to report your switching point.
- 2. Only one question in one of the eight blocks will be randomly selected for payment.
- 3. You will be paid on the designated date through VENMO.

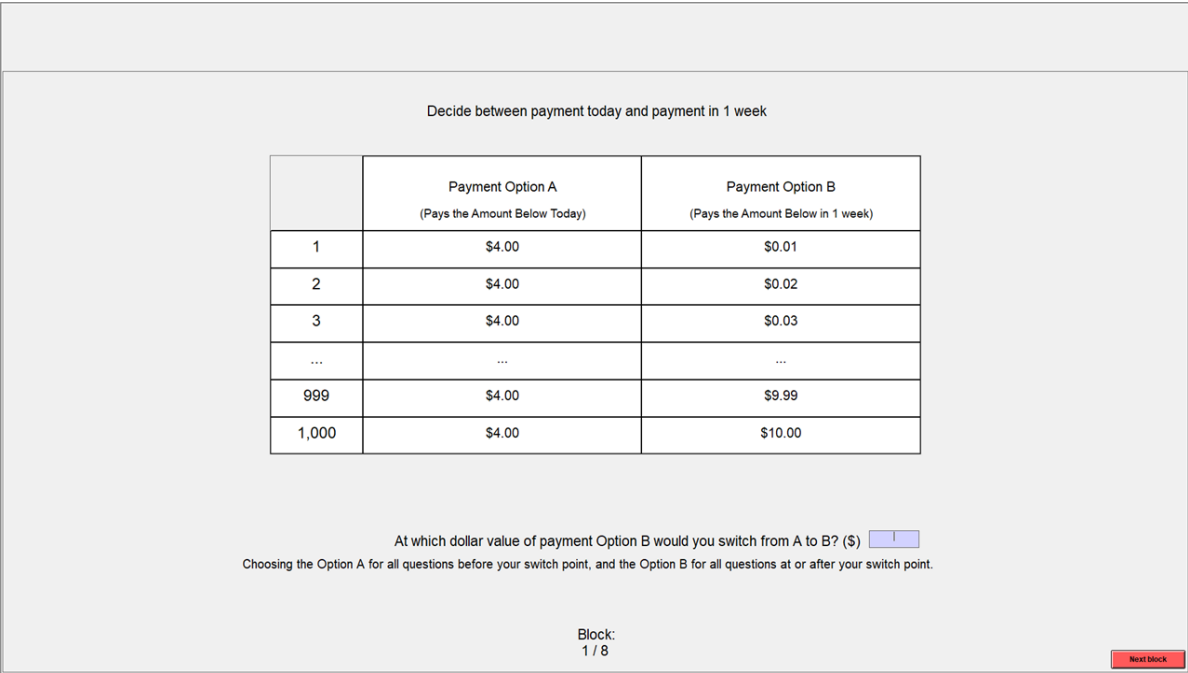


Figure 24: Elicitation Task Screen-shot Block 1